

Belcore

384

GENERIC REQUIREMENTS
GR-253-CORE
ISSUE 2, DECEMBER 1995
WITH REVISION 2, JANUARY 1999

BEST AVAILABLE COPY

Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria

(A Module of TSGR, FR-440)

Belcore

GENERIC REQUIREMENTS

GR-253-CORE

ISSUE 2, DECEMBER 1995

WITH REVISION 2, JANUARY 1999

Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria

(A Module of TSGR, FR-440)

This document includes all of GR-253-CORE, Issue 2, December 1995, including Revision 1, December 1997, and Revision 2, January 1999. GR-253-CORE, Issue 2 completely replaced GR-253-CORE, Issue 1.

For ordering information, see the References section of this document.

This document may not be reproduced without the express written permission of Bellcore and any reproduction without written authorization is an infringement of Bellcore's copyright.

Copyright © 1999 Bellcore.

All rights reserved.

Project funding year: 1998.

Trademark Acknowledgments

CLEI is a trademark of Bellcore

CLLI is a trademark of Bellcore

GENERIC REQUIREMENTS NOTICE OF DISCLAIMER

This Generic Requirements document (GR) is published by Bellcore to inform the industry of Bellcore's view of proposed generic requirements on the Synchronous Optical Network (SONET). These generic requirements are subject to review and change, and superseding generic requirements regarding this subject may differ from this document. Bellcore reserves the right to revise this document for any reason (consistent with applicable provisions of the Telecommunications Act of 1996 and applicable FCC rules).

Bellcore specifically advises the reader that this GR does not directly or indirectly address any Year-2000 ("Y2K") issues that might be raised by the services, systems, equipment, specifications, descriptions, or interfaces addressed or referred to herein. As an example, and not a limitation, neither this GR nor Bellcore is directly or indirectly assessing or determining whether specific services, systems, or equipment, individually or together, in their current form, or as they may be implemented, modified, or augmented in the future, will accurately process dates and date-related data within or between the twentieth and twenty-first centuries, in either direction, including elapsed time, time difference, and/or leap year calculations.

BELLCORE, AND FUNDING PARTICIPANTS IDENTIFIED IN THE PREFACE, MAKE NO REPRESENTATION OR WARRANTY, EXPRESSED OR IMPLIED, WITH RESPECT TO THE SUFFICIENCY, ACCURACY, OR UTILITY OF ANY INFORMATION OR OPINION CONTAINED HEREIN.

BELLCORE, AND FUNDING PARTICIPANTS, EXPRESSLY ADVISE THAT ANY USE OF OR RELIANCE UPON SAID INFORMATION OR OPINION IS AT THE RISK OF THE USER AND THAT NEITHER BELLCORE, NOR ANY FUNDING PARTICIPANT, SHALL BE LIABLE FOR ANY DAMAGE OR INJURY INCURRED BY ANY PERSON ARISING OUT OF THE SUFFICIENCY, ACCURACY, OR UTILITY OF ANY INFORMATION OR OPINION CONTAINED HEREIN.

LOCAL CONDITIONS MAY GIVE RISE TO A NEED FOR ADDITIONAL PROFESSIONAL INVESTIGATIONS, MODIFICATIONS, OR SAFEGUARDS TO MEET SITE, EQUIPMENT, ENVIRONMENTAL SAFETY OR COMPANY-SPECIFIC REQUIREMENTS. IN NO EVENT IS THIS INFORMATION INTENDED TO REPLACE FEDERAL, STATE, LOCAL, OR OTHER APPLICABLE CODES, LAWS, OR REGULATIONS. SPECIFIC APPLICATIONS WILL CONTAIN VARIABLES UNKNOWN TO OR BEYOND THE CONTROL OF BELLCORE. AS A RESULT, BELLCORE CANNOT WARRANT THAT THE APPLICATION OF THIS INFORMATION WILL PRODUCE THE TECHNICAL RESULT OR SAFETY ORIGINALLY INTENDED.

This GR is not to be construed as a suggestion to anyone to modify or change any of its products or services, nor does this GR represent any commitment by anyone, including, but

not limited to, Bellcore or any funder (see Preface) of this Bellcore GR to purchase, manufacture, or sell, any product with the described characteristics.

Readers are specifically advised that any entity may have needs, specifications, or requirements different from the generic descriptions herein. Therefore, anyone wishing to know any entity's needs, specifications, or requirements should communicate directly with that entity.

Nothing contained herein shall be construed as conferring by implication, estoppel, or otherwise any license or right under any patent, whether or not the use of any information herein necessarily employs an invention of any existing or later issued patent.

Nothing contained herein is intended as a recommendation of any product to anyone.

TSGR Contents

Vol	Section	Module
1	Common Criteria	Common Requirements, GR-499
		NEBS, GR-63
		Supplier Documentation, GR-454
		Supplier-Provided Training, GR-839
		Supplier Support, TR-NWT-000840
		EMC and Electrical Safety, GR-1089
		Assuring Corrosion Resistance of Outside Equipment, GR-2836
2	DLC Systems, UDC, and Basic Exchange Radio Systems	Digital Loop Carrier Systems, TR-TSY-000057
		Universal Digital Channel (UDC), TR-TSY-000398
		Basic Exchange Radio Systems, TR-NWT-000911
3	IDLC System	IDLC System Criteria, GR-303
4	ISDN Transport, Interface, and Related Requirements	ISDN Basic Access DSL, TR-NWT-000393
		ISDN Basic Access Transport, TR-NWT-000397
		ISDN PRA Transport, TR-TSY-000754
5	Digital Radio Systems	Microwave Digital Radio, TR-TSY-000752
6	SONET Common Criteria	SONET Transport Systems: Common Generic Criteria, GR-253
		SONET OC-192 Criteria, GR-1377
		SONET Private Line Service Interface Criteria, GR-1365
		SONET Inter-Carrier Interface Criteria, GR-1374
7	SONET Network Element Criteria	SONET Add/Drop Multiplex, GR-496
		Wideband and Broadband Digital Cross-Connect, TR-NWT-000233
		SONET Digital Switch Trunk Interface, TR-TSY-000782
		SONET RGTR Equipment Criteria, TR-NWT-000917
		SONET Dual-Fed UPSR Criteria, GR-1400
		SONET BLSR Criteria, GR-1230
		Self-Healing Ring Functionality in Digital Cross-Connect Systems, GR-1375
		ATM Virtual Path Functionality in SONET Rings, GR-2837
		SONET ATM Virtual Path Digital Cross-Connect Systems - Generic Criteria, GR-2891

NOTE:

This document is a module of the *Transport Systems Generic Requirements (TSGR)*, FR-440. Modules or the entire TSGR can be ordered from Bellcore's on-line catalog or by contacting Bellcore.

To Contact Bellcore

Bellcore Customer Service
8 Corporate Place, Room 3A-184
Piscataway, New Jersey 08854-4156
1-800-521-CORE (2673) (USA and Canada)
1 (732) 699-5800 (all others)
1 (732) 336-2559 (FAX)

To Order Documents On-Line

- Perform the following steps to order from the on-line catalog:

1. Enter the URL line: *telecom-info.bellcore.com*
2. Click on the *Search* button located on top
3. In the *Keywords* field, enter the document number (or keywords), then click on *Submit Search*.

or . . .

1. Enter the above URL line
2. Click on the *Browse* button located on top, then click on the subject of interest.

- Bellcore employees should perform the following steps:

1. Access the Bellcore Internal Home Page
2. Click on the *Marketwise* button located on top
3. Click on *DOCS - Bellcore Product Catalog*
4. Enter data in one or more of the fields (e.g., enter a document number in the *Product Number* field), then follow the instructions to search the on-line catalog.

Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria

Contents

Preface	Preface-1
1. Introduction	1-1
1.1 Requirements Terminology	1-1
1.2 Requirement Labeling Conventions	1-2
1.2.1 Numbering of Requirement and Related Objects	1-2
1.2.2 Requirement, Conditional Requirement, and Objective Object Identification	1-2
1.3 Revision History	1-3
1.3.1 Changes Between GR-253-CORE Issues 1 and 2	1-4
1.3.2 Changes Between GR-253-CORE Issue 2 and Revision 1 to Issue 2	1-8
1.3.3 Changes In Revision 2 to GR-253-CORE Issue 2	1-13
2. Network Compatibility	2-1
2.1 Network Element (NE) Interfaces	2-2
2.1.1 Digital Signal Cross-Connect (DSX) Interface	2-2
2.1.1.1 Electrical Cable Distance	2-3
2.1.1.2 Maintenance Signal Compatibility	2-3
2.1.2 Interface to Fiber Distributing Frame and Optical DSXs	2-3
2.1.3 Operations Systems Interface	2-4
2.1.4 Synchronization Interface	2-4
2.1.5 Power	2-4
2.2 End-to-End Performance Criteria	2-5
2.2.1 Availability and Reliability	2-5
2.2.2 Protection Switching Performance	2-6
2.2.3 Error Performance	2-6
2.2.4 Jitter	2-6
2.2.5 Transmission Delay	2-6
3. Rates and Formats	3-1
3.1 Synchronous Hierarchical Rates	3-1
3.2 Transport Format	3-2
3.2.1 Frame Structure of the STS-1	3-3
3.2.1.1 Transport Overhead	3-4
3.2.1.2 STS-1 Envelope Capacity and Synchronous Payload Envelope (SPE)	3-4
3.2.2 Frame Structure of the STS-N	3-7
3.2.3 STS Concatenation	3-7
3.2.4 Virtual Tributary (VT) Structure	3-10
3.3 Layered Overhead and Transport Functions	3-24

3.3.1	SONET Interface Layers.....	3-24
3.3.1.1	Physical Layer	3-24
3.3.1.2	Section Layer	3-25
3.3.1.3	Line Layer.....	3-25
3.3.1.4	Path Layer.....	3-25
3.3.1.5	Interaction of the Layers.....	3-26
3.3.2	STS-1 Overhead Descriptions.....	3-27
3.3.2.1	Section Overhead.....	3-29
3.3.2.2	Line Overhead	3-30
3.3.2.3	STS Path Overhead (STS POH)	3-32
3.3.3	VT Path Overhead.....	3-38
3.4	Payload Mapping	3-41
3.4.1	Sub-STS-1 Mappings	3-41
3.4.1.1	Byte-Synchronous Mapping for DS1	3-42
3.4.1.2	Asynchronous Mapping for DS1	3-49
3.4.1.3	Asynchronous Mapping for DS1C	3-51
3.4.1.4	Asynchronous Mapping for DS2	3-52
3.4.2	STS-1 Mappings.....	3-55
3.4.2.1	Asynchronous Mapping for DS3	3-55
3.4.2.2	Asynchronous Transfer Mode (ATM) Mapping for B-ISDN Applications.....	3-56
3.4.2.3	HDL-Over-SONET Mapping	3-57
3.4.3	Super Rate Payloads.....	3-58
3.4.3.1	Mappings into an STS-3c SPE.....	3-58
3.4.3.2	Mappings into STS-12c SPEs	3-65
3.5	Payload Pointers.....	3-66
3.5.1	STS Payload Pointer	3-66
3.5.1.1	Pointer Value	3-67
3.5.1.2	STS Frequency Justification	3-68
3.5.1.3	New Data Flag (NDF)	3-72
3.5.1.4	Concatenation Indicator.....	3-72
3.5.1.5	STS Payload Pointer Generation Rules	3-73
3.5.1.6	STS Payload Pointer Interpretation	3-74
3.5.2	VT Payload Pointer.....	3-75
3.5.2.1	VT Pointer Value.....	3-76
3.5.2.2	VT Frequency Justification.....	3-77
3.5.2.3	VT Size Indicator.....	3-77
3.5.2.4	New Data Flag (NDF)	3-78
3.5.2.5	VT Payload Pointer Generation Rules.....	3-79
3.5.2.6	VT Payload Pointer Interpretation Rules.....	3-80
4.	Physical Layer.....	4-1
4.1	Physical Layer Classifications	4-1
4.2	Optical Parameter Definitions and Interface Requirements.....	4-2
4.2.1	General	4-2

4.2.2	Optical Line Coding.....	4-4
4.2.3	Reflections	4-4
4.2.4	Transmitter	4-6
4.2.4.1	Spectral Characteristics	4-7
4.2.4.2	Coupled Transmit Power	4-7
4.2.4.3	Extinction Ratio	4-8
4.2.4.4	Mask of the Eye Diagram	4-8
4.2.5	Receiver	4-12
4.2.6	Optical Path.....	4-12
4.3	Engineering of a Single-Mode Fiber Optic Transmission System.....	4-23
4.3.1	Terminal Equipment Transmission Design Information	4-23
4.3.1.1	General System Information.....	4-24
4.3.1.2	Transmitter Information.....	4-24
4.3.1.3	Receiver Information.....	4-25
4.3.1.4	Attenuators.....	4-27
4.3.1.5	Wavelength Division Multiplex (WDM) Device	4-27
4.3.1.6	Safety Margin	4-28
4.3.1.7	Connectors	4-28
4.3.1.8	Station Cable.....	4-29
4.3.2	Cable Transmission Design Information	4-30
4.3.2.1	Parameters for a Specific Application	4-30
4.3.2.2	Global Fiber Parameters	4-33
4.3.3	Fiber Optic System Transmission Design and Analysis	4-37
4.3.3.1	Design Approach	4-37
4.3.3.2	Loss Budget Constraint.....	4-39
4.3.3.3	Dispersion Limited Length.....	4-41
4.3.3.4	Design and Analysis Methodology.....	4-41
4.4	Electrical Interface Specifications.....	4-42
5.	Network Element Architectural Features.....	5-1
5.1	Multiplex Procedures	5-1
5.1.1	Interleaving	5-1
5.1.2	Concatenation.....	5-6
5.1.3	Scrambling	5-6
5.1.4	An Example of STS-1 and OC-N Signal Composition	5-7
5.2	Overhead Function Usage	5-10
5.2.1	Generating and Processing Overhead	5-10
5.2.2	Orderwire (OW).....	5-13
5.2.2.1	OW Access	5-14
5.2.2.2	OW System Communication	5-15
5.2.2.3	OW Operations	5-16
5.2.3	User Channels	5-16
5.3	Automatic Protection Switching	5-16
5.3.1	Protection Switching Boundaries.....	5-17
5.3.2	Linear APS Architectures	5-17

5.3.2.1	1+1 Architecture	5-18
5.3.2.2	1:n Architecture	5-19
5.3.2.3	1:1 Case of the 1:n Architecture	5-20
5.3.3	Switch Initiation and Completion Criteria	5-21
5.3.3.1	Switch Initiation Criteria	5-21
5.3.3.2	Switch Initiation Time	5-21
5.3.3.3	Switch Completion Time	5-23
5.3.4	Restoral and Clearing of SD and SF Conditions.....	5-24
5.3.5	APS Channel Protocol	5-28
5.3.5.1	K1 Byte	5-28
5.3.5.2	K2 Byte	5-34
5.3.5.3	Control of the Bridge	5-36
5.3.5.4	Control of the Selector	5-38
5.3.5.5	Transmission and Acceptance of Bytes K1 and K2	5-39
5.3.6	Linear APS Commands	5-41
5.3.6.1	Switch Commands	5-42
5.3.6.2	Control Commands	5-43
5.3.7	Switch Operation.....	5-44
5.3.7.1	1:n Architecture	5-44
5.3.7.2	1+1 Architecture	5-49
5.4	Network Synchronization.....	5-50
5.4.1	SONET NE Clock Applications	5-50
5.4.1.1	Physical Interface to Synchronization Network	5-52
5.4.1.2	TDEV and MTIE Measurements.....	5-52
5.4.2	Synchronization Status Messages	5-52
5.4.3	SONET Timing Modes	5-56
5.4.3.1	External Timing	5-57
5.4.3.2	Line-Timing.....	5-58
5.4.3.3	Loop-Timing.....	5-59
5.4.3.4	Through-Timing	5-59
5.4.4	SONET Internal Clock.....	5-60
5.4.4.1	Stratum Clocks for SONET Applications.....	5-61
5.4.4.2	SONET Minimum Clock Applications	5-61
5.4.4.3	All SONET Clocks	5-67
5.4.5	Timing Distribution.....	5-73
5.4.5.1	Timing Distribution on Derived DS1 Signals	5-74
5.4.5.2	Synchronization Status Messages for Derived DS1 Signals.....	5-81
5.4.5.3	Timing Distribution on Traffic-Carrying DS1 Payload Signals.....	5-83
5.4.6	SONET Timing Reference Switching.....	5-84
5.4.6.1	Timing Reference Failure Conditions	5-87
5.4.6.2	Performance During Timing Reference Switching	5-89

5.4.6.3	Revertive and Nonrevertive Timing Reference Switching	5-89
5.4.6.4	Synchronization Status Messages and Timing Reference Switching	5-89
5.4.7	Message Validation and Generation	5-91
5.4.7.1	Message Validation	5-91
5.4.7.2	Message Reaction	5-93
5.4.7.3	Message Generation	5-93
5.5	Framing For SONET Signals	5-97
5.6	Jitter	5-98
5.6.1	Network Interface Jitter Criteria	5-99
5.6.2	SONET NE Jitter Criteria	5-100
5.6.2.1	Jitter Transfer	5-101
5.6.2.2	Jitter Tolerance	5-105
5.6.2.3	Jitter Generation	5-107
5.6.2.4	Jitter Enhancement	5-115
5.7	Phase Variations on Payload Signals	5-119
5.7.1	Mapping Phase Variations	5-119
5.7.2	Pointer Adjustment Phase Variations	5-121
5.7.2.1	Single Pointer Adjustments	5-122
5.7.2.2	Pointer Adjustment Bursts	5-124
5.7.2.3	Periodic Pointer Adjustments	5-126
6.	SONET Network Element Operations Criteria	6-1
6.1	Memory Administration	6-1
6.1.1	Memory Administration Data	6-1
6.1.2	Data Manipulation	6-2
6.1.3	Administration of Operations Communications Information	6-4
6.1.4	Regenerators	6-4
6.1.5	Memory Backup and Restoration	6-5
6.1.6	System Administration and Security	6-6
6.1.6.1	NE Security Mechanism	6-7
6.1.6.2	DCC Security	6-10
6.1.7	Software Generics	6-12
6.1.8	Self-Inventory	6-13
6.2	Maintenance	6-13
6.2.1	Alarm Surveillance	6-14
6.2.1.1	Directly Detected Defects and Failures	6-16
6.2.1.2	Alarm Indication Signal	6-42
6.2.1.3	Remote Defect Indication (RDI) and Remote Failure Indication (RFI)	6-51
6.2.1.4	Payload Defect Indication (PDI)	6-63
6.2.1.5	Maintenance Signals for Other Mappings	6-66
6.2.1.6	Trunk Conditioning	6-66
6.2.1.7	Alarm-Related Events	6-68

6.2.1.8	Control of Alarm Processing	6-83
6.2.2	Performance Monitoring	6-92
6.2.2.1	General Accumulation and Thresholding Criteria	6-98
6.2.2.2	Physical Layer PM	6-104
6.2.2.3	Section Layer PM	6-107
6.2.2.4	Line Layer PM	6-109
6.2.2.5	STS Path Layer PM	6-112
6.2.2.6	VT Path Layer PM	6-117
6.2.2.7	Monitoring at DS _n Interfaces	6-121
6.2.2.8	PM During Troubles	6-121
6.2.2.9	Intermediate-Path PM	6-129
6.2.3	Testing Process	6-135
6.2.3.1	Test Access	6-135
6.2.3.2	Diagnostics	6-137
6.2.3.3	Loopbacks	6-144
6.2.4	Control Features	6-149
7.	Other Generic Criteria	7-1
7.1	Physical and Environmental Criteria	7-1
7.1.1	Operational Environment for Equipment	7-1
7.1.2	Outside Plant Cable	7-2
7.2	Equipment Design	7-2
7.3	Documentation and Training	7-2
7.4	Safety	7-3
7.4.1	Station Equipment Safety	7-3
7.4.2	Fiber Optic Cable Safety	7-4
7.5	Quality and Reliability	7-4
7.5.1	Network Equipment Reliability	7-4
7.5.2	Fiber Optic Cable Quality and Reliability	7-4
7.5.3	Component Reliability Assurance	7-4
7.6	Other Requirements and Objectives	7-5
7.6.1	Human Factors	7-5
7.6.2	Technical Analysis	7-5
8.	SONET Operations Communications	8-1
8.1	SONET Operations Communications Architecture	8-2
8.1.1	Architecture Overview	8-2
8.1.2	Gateway NE Requirements	8-6
8.1.3	Intermediate NE Requirements	8-8
8.1.4	End NE Requirements	8-8
8.1.5	Mediation Device (MD)	8-8
8.2	Communication Types	8-9
8.2.1	OS/NE Communications	8-9
8.2.2	MD/NE Communications	8-10
8.2.3	NE/NE Communications	8-10

8.2.4	Craftsperson/NE Communications	8-11
8.3	SONET Operations Communications Interface	8-11
8.3.1	Physical Layer	8-13
8.3.1.1	OS/NE	8-13
8.3.1.2	NE/NE - LAN	8-13
8.3.1.3	NE/NE - DCC	8-13
8.3.2	Data Link Layer	8-14
8.3.2.1	OS/NE	8-14
8.3.2.2	NE/NE - LAN	8-15
8.3.2.3	NE/NE - DCC	8-15
8.3.3	Network Layer	8-16
8.3.3.1	OS/NE	8-16
8.3.3.2	NE/NE - LAN and DCC	8-17
8.3.4	Transport Layer - OS/NE and NE/NE	8-20
8.3.5	Session Layer - OS/NE and NE/NE	8-21
8.3.6	Presentation Layer - OS/NE and NE/NE	8-22
8.3.7	Application Layer - OS/NE and NE/NE	8-22
8.3.7.1	ACSE	8-22
8.3.7.2	ROSE/CMISE	8-22
8.3.7.3	FTAM	8-23
8.3.7.4	Name/Address Translation Services	8-23
8.3.7.5	TL1	8-23
8.4	Interworking between OSs and SONET NEs	8-25
8.4.1	TL1/X.25 [OS]-TL1/OSI [SONET]	8-26
8.4.1.1	Determine Destination NSAP	8-26
8.4.1.2	Directing Autonomous Messages	8-28
8.4.1.3	Establishing Connections	8-29
8.4.1.4	SONET LAN Interworking	8-30
8.4.1.5	SONET DCC Interworking	8-31
8.4.1.6	SONET LAN and DCC Interworking	8-31
8.4.2	TL1/X.25 [OS]-CMISE/OSI [SONET]	8-32
8.4.3	CMISE/OSI [OS]-CMISE/OSI [SONET]	8-32
8.4.4	CMISE (or TL1)/OSI [SONET]-CMISE (or TL1)/OSI [SONET]	8-33
8.5	SONET Operations Communications Routing	8-33
8.5.1	Routing Overview	8-33
8.5.2	ES-IS Requirements	8-35
8.5.3	IS-IS Requirements	8-36
8.6	Craftsperson/NE Interfaces	8-37
8.6.1	Craftsperson/WS Interface	8-38
8.6.2	WS/NE Interface	8-39
8.7	TARP	8-39
8.7.1	Network Layer Protocol to Support TARP	8-40
8.7.2	TARP PDU Specification	8-40
8.7.2.1	TARP Lifetime (tar-lif)	8-41

8.7.2.2	TARP Sequence Number (tar-seq)	8-41
8.7.2.3	Protocol Address Type (tar-pro)	8-41
8.7.2.4	URC and TARP Type Code (tar-tcd)	8-41
8.7.2.5	TID Target Length (tar-tln)	8-42
8.7.2.6	TID Originator Length (tar-oln)	8-42
8.7.2.7	Protocol Address Length (tar-pln)	8-42
8.7.2.8	TID of Target (tar-ttg)	8-43
8.7.2.9	TID of Originator (tar-tor)	8-43
8.7.2.10	Protocol Address of Originator (tar-por)	8-43
8.7.3	TARP Data Cache (TDC)	8-43
8.7.4	NE Applications That Use the TARP Processor	8-43
8.7.4.1	Find NET That Matches TID	8-44
8.7.4.2	Find TID That Matches NET	8-44
8.7.4.3	Send Notification of TID or Protocol Address Change	8-45
8.7.5	TARP PDU Processing	8-45
8.7.5.1	Origination of a TARP Type 1 PDU	8-46
8.7.5.2	Origination of a TARP Type 2 PDU	8-46
8.7.5.3	Origination of a TARP Type 3 PDU	8-46
8.7.5.4	Origination of a TARP Type 4 PDU	8-47
8.7.5.5	Origination of a TARP Type 5 PDU	8-47
8.7.5.6	Receipt of a TARP PDU	8-47
8.7.5.7	Loop Detection Procedure (performed by ISs)	8-49
8.7.5.8	Propagation Procedure (performed by ISs)	8-49
8.7.6	Management of the TARP Processor	8-50
8.7.7	TARP Echo Function (TEF)	8-51
8.7.8	Manual TARP Adjacencies	8-52
8.7.9	TARP Example	8-52
8.7.10	TARP Pseudocode	8-54
Appendix A: Requirement-Object List		A-1
Appendix B: Fiber Optic Transmission System Design Worksheets		B-1
Appendix C: SONET Operations Communications		
Lower Layers Protocol Profile		C-1
C.1	Introduction	C-1
C.1.1	Source Documents	C-1
C.1.1.1	Base Standards	C-1
C.1.1.2	PICS Proforma	C-2
C.1.1.3	International Standardized Profiles	C-3
C.1.1.4	American National Standards for Telecommunications	C-3
C.1.1.5	Bellcore Requirements	C-3
C.1.2	Notations Used in the SONET Lower Layer Profile	C-3
C.1.2.1	Status Symbols	C-3
C.1.2.2	Profile Symbols	C-3
C.1.2.3	Support Symbols	C-5

	C.1.2.4	References.....	C-5
C.2		LAPD SONET Protocol Profile.....	C-6
	C.2.1	Notations.....	C-6
		C.2.1.1 Abbreviations.....	C-6
		C.2.1.2 Status Symbols.....	C-6
		C.2.1.3 Additional Symbols.....	C-6
	C.2.2	Protocol Capabilities (PC).....	C-7
	C.2.3	Frames - Protocol Data Units (FR).....	C-9
	C.2.4	System Parameters (SP).....	C-10
C.3		LLC Protocol Profile.....	C-11
	C.3.1	Abbreviations and Special Symbols.....	C-11
		C.3.1.1 Status Symbols.....	C-11
		C.3.1.2 Item References.....	C-11
	C.3.2	Claimed Conformance to ISO 8802-2:1989/Amd.1, Amd.2 and Amd.4.....	C-12
	C.3.3	LLC Type 1 Operation - Unacknowledged Connectionless Mode.....	C-12
		C.3.3.1 LLC Type 1 - Supported PDU Types.....	C-12
		C.3.3.2 LLC Type 1 - Supported Parameters in PDUs on Transmission.....	C-13
		C.3.3.3 LLC Type 1 - Supported Parameters in PDUs on Receipt.....	C-14
		C.3.3.4 LLC Type 1 - Miscellaneous.....	C-14
C.4		ISO 8473 Protocol Profile.....	C-16
	C.4.1	Notations.....	C-16
		C.4.1.1 Status Symbols.....	C-16
		C.4.1.2 Additional Symbols.....	C-16
	C.4.2	Major Capabilities.....	C-16
	C.4.3	End Systems.....	C-17
		C.4.3.1 Applicability.....	C-17
		C.4.3.2 Supported Functions.....	C-17
		C.4.3.3 Supported PDUs.....	C-18
		C.4.3.4 Supported Parameters.....	C-18
		C.4.3.5 Timers.....	C-19
	C.4.4	Intermediate Systems.....	C-20
		C.4.4.1 Applicability.....	C-20
		C.4.4.2 Supported Functions.....	C-20
		C.4.4.3 Supported PDUs.....	C-21
		C.4.4.4 Supported Parameters.....	C-21
		C.4.4.5 Timer and Parameter Values.....	C-22
C.5		ISO 9542 SONET Protocol Profile.....	C-23
	C.5.1	Notations.....	C-23
		C.5.1.1 Status Symbols.....	C-23
		C.5.1.2 Other Symbols.....	C-23
	C.5.2	PICS Proforma: ISO 9542(1988) - End System.....	C-23

C.5.2.1	Supported Functions	C-24
C.5.2.2	Supported PDUs	C-24
C.5.2.3	Supported Parameters	C-25
C.5.2.4	Supported Parameter Ranges	C-25
C.5.3	PICS Proforma: ISO 9542(1988) – Intermediate System	C-26
C.5.3.1	Supported Functions	C-26
C.5.3.2	Supported PDUs	C-26
C.5.3.3	Supported Parameters	C-27
C.5.4	Supported Parameter Ranges	C-27
C.6	ISO 10589 Protocol Profile	C-28
C.6.1	Notations (Status Symbols)	C-28
C.6.2	Protocol Summary: ISO 10589 General	C-28
C.6.2.1	System Environment: General	C-29
C.6.2.2	Subnetwork Dependent Functions: General	C-29
C.6.2.3	Update Process: General	C-30
C.6.2.4	Decision Process: General	C-31
C.6.2.5	Forward/Receive Process: General	C-31
C.6.3	Protocol Summary: ISO 10589 Level 1 Specific Functions	C-32
C.6.3.1	Level 1 Subnetwork Dependent Functions	C-32
C.6.3.2	Level 1 Update Process	C-32
C.6.3.3	Level 1 Decision Process	C-32
C.6.4	Protocol Summary: ISO 10589 Level 2 Specific Functions	C-33
C.6.4.1	Level 2 Subnetwork Dependent Functions	C-33
C.6.4.2	Level 2 Update Process	C-34
C.6.4.3	Level 2 Decision Process	C-34
C.6.4.4	Level 2 Forward/Receive Process	C-34
C.7	ISO 8073 Protocol Profile	C-35
C.7.1	Notations	C-35
C.7.1.1	Status Symbols	C-35
C.7.2	Protocol Implementation for TP4/CLNS (C4L:)	C-35
C.7.2.1	Annex B – NCMS	C-35
C.7.2.2	Classes Implemented	C-35
C.7.3	Initiator/Responder Capability for Protocol Classes 0-4	C-36
C.7.4	Supported Functions	C-36
C.7.4.1	Supported Functions for Class 4 (C4L:)	C-36
C.7.5	Supported TPDU's	C-38
C.7.6	Supported Parameters of Issued TPDU's	C-39
C.7.6.1	Parameter Values for CR TPDU (C4L:)	C-39
C.7.6.2	Supported Parameters for Class 4 TPDU's (C4L:)	C-39
C.7.7	Supported Parameters for Received TPDU's	C-41
C.7.8	User Data in Issued TPDU's	C-41
C.7.8.1	Class 4 (C4L:)	C-41
C.7.9	User Data in Received TPDU's	C-42
C.7.10	Negotiation	C-42

C.7.10.1	Class Negotiation - Initiator.....	C-42
C.7.10.2	TPDU Size Negotiation.....	C-43
C.7.10.3	Use of Extended Format.....	C-43
C.7.10.4	Expedited Data Transport Service.....	C-44
C.7.10.5	Non-use of Checksum (C4L AND T4F29::).....	C-44
C.7.10.6	Use of Selective Acknowledgment (See note 20).....	C-44
C.7.10.7	Use of Request of Acknowledgment (See note 21).....	C-44
C.7.11	Error Handling.....	C-45
C.7.11.1	Action on Receipt of a Protocol Error.....	C-45
C.7.11.2	Actions on Receipt of an Invalid or Undefined Parameter in a CR TPDU.....	C-45
C.7.11.3	Actions on Receipt of an Invalid or Undefined Parameter in a TPDU other than a CR TPDU.....	C-46
C.7.12	Timers and Protocol Parameters.....	C-46
C.8	TARP Protocol Implementation Conformance Statement.....	C-47
C.8.1	Major Function.....	C-47
C.8.2	Supported PDUs.....	C-48
C.8.3	Protocol Specifications.....	C-48
C.8.3.1	TARP PDU CLNP Specifications.....	C-48
C.8.3.2	TARP PDU Specifications.....	C-49
C.8.3.3	Protocol Timer Specifications.....	C-49
C.8.4	Major Capabilities.....	C-50
C.8.5	TARP Processor Management.....	C-51
C.8.5.1	LDB Entry Timer Parameters.....	C-52
C.8.5.2	LDB Flush Timer Parameters.....	C-52
C.8.5.3	Provisionable TARP PDU Fields.....	C-52
Appendix D:	SONET Operations Communications.....	D-1
Upper Layers	Protocol Profile.....	D-1
D.1	Introduction.....	D-1
D.2	Source Documents.....	D-1
D.2.1	Base Standards.....	D-1
D.2.2	PICS Proforma.....	D-1
D.2.3	International Standardized Profiles.....	D-2
D.2.4	Bellcore Requirements.....	D-2
D.3	Goals of SONET Upper Layers Profile.....	D-2
D.4	Structure of SONET Upper Layers Profile.....	D-3
D.5	Notations Used in the SONET Upper Layers Profile.....	D-3
D.5.1	Abbreviations.....	D-3
D.5.2	Status Column.....	D-4
D.5.3	Profile Column.....	D-4
D.5.4	Support Column.....	D-4
D.5.5	PICS Numbers.....	D-4
D.6	SONET Upper Layers Profile: ACSE.....	D-6
D.6.1	Additions Beyond Existing ISP Requirements.....	D-6

D.6.2	Profile Tables	D-6
D.7	SONET Upper Layers Profile: Presentation Layer	D-11
D.7.1	Additions Beyond Existing ISP Requirements	D-11
D.7.2	Profile Tables	D-11
D.8	SONET Upper Layers Profile: Session Layer	D-18
D.8.1	Additions Beyond Existing ISP Requirements	D-18
D.8.2	Profile Tables	D-18
References	References-1
Glossary	Glossary-1
Requirement-Object Index	ROI-1

List of Figures

Figure 2-1.	Simplified Diagram Depicting SONET Section, Line, and Path Definitions	2-1
Figure 2-2.	Diagram Illustrating SONET Section, Line, and Path Definitions.....	2-2
Figure 3-1.	STS-1 Frame	3-3
Figure 3-2.	Bit Position Numbering	3-4
Figure 3-3.	STS-1 Synchronous Payload Envelope	3-5
Figure 3-4.	STS-1 SPE with STS-1 POH and STS-1 Payload Capacity Illustrated	3-6
Figure 3-5.	STS-1 SPE in Interior of STS-1 Frames	3-6
Figure 3-6.	STS-N Frame	3-7
Figure 3-7.	STS-Nc SPE.....	3-9
Figure 3-8.	Transport Overhead Assignment, OC-3 Carrying an STS-3c SPE.....	3-9
Figure 3-9.	VT Sizes	3-11
Figure 3-10.	VT Structured STS-1 SPE: All VT1.5s	3-12
Figure 3-11.	VT1.5 Locations	3-13
Figure 3-12.	VT Structured STS-1 SPE: All VT2s	3-14
Figure 3-13.	VT2 Locations	3-15
Figure 3-14.	VT Structured STS-1 SPE: All VT3s	3-16
Figure 3-15.	VT3 Locations	3-17
Figure 3-16.	VT Structured STS-1 SPE: All VT6s	3-18
Figure 3-17.	VT6 Locations	3-19
Figure 3-18.	Example of VT Structured STS-1 SPE	3-20
Figure 3-19.	Correspondence Between Labels and Numbers for the Example in Figure 3-18	3-21
Figure 3-20.	VT Superframe and Envelope Capacity	3-22
Figure 3-21.	VT SPE and Payload Capacity	3-23
Figure 3-22.	Optical Interface Layers	3-27
Figure 3-23.	Transport and Path Overhead Byte Designations.....	3-28
Figure 3-24.	STS Path Status Byte (G1)	3-37
Figure 3-25.	VT Path Overhead Byte (V5)	3-38
Figure 3-26.	H4 Byte Coding Sequence for VT-Structured STS-1 SPEs	3-41
Figure 3-27.	Byte-Synchronous Mapping for DS1 Payload	3-43
Figure 3-28.	Byte-Synchronous DS1 Signaling and Framing Bit Assignments	3-44
Figure 3-29.	Asynchronous Mapping for DS1 Payload	3-50
Figure 3-30.	Asynchronous Mapping for DS1C Payload	3-52
Figure 3-31.	Asynchronous Mapping for DS2 Payload	3-54
Figure 3-32.	Asynchronous Mapping for DS3 Payload	3-56
Figure 3-33.	Asynchronous Mapping for DS4NA Payload	3-60
Figure 3-34.	Asynchronous Mapping for FDDI.....	3-62
Figure 3-35.	DQDB Mapping into an STS-3c SPE.....	3-64
Figure 3-36.	Bit Allocation for H4 Byte in DQDB Mapping	3-64
Figure 3-37.	STS Payload Pointer (H1, H2) Coding.....	3-67

Figure 3-38.	STS Pointer Offset Numbering.....	3-68
Figure 3-39.	Positive STS Pointer Adjustment Operation (Increment)	3-70
Figure 3-40.	Negative STS Pointer Adjustment Operation (Decrement)	3-71
Figure 3-41.	VT Payload Pointer (V1, V2) Coding	3-75
Figure 3-42.	VT Pointer Offsets	3-76
Figure 4-1.	Optical System Interfaces (Points S and R).....	4-3
Figure 4-2.	SONET Eye Diagram Mask (OC-1 to OC-12)	4-10
Figure 4-3.	SONET Eye Diagram Mask (OC-48)	4-11
Figure 4-4.	Cable Configuration for Cabled Fiber Cutoff Wavelength Measurement	4-31
Figure 4-5.	Wavelength Dependent Attenuation Characteristics	4-32
Figure 4-6.	Global Fiber Attenuation Characteristics	4-34
Figure 4-7.	Global Fiber Dispersion Characteristics	4-36
Figure 4-8.	Engineering of a Fiber Optic Regenerator Section.....	4-38
Figure 4-9.	Regenerator Section, Sources of Loss	4-40
Figure 4-10.	STS-1 Electrical Interface Pulse Mask	4-45
Figure 4-11.	STS-1 Electrical Interface Eye Diagram Mask.....	4-46
Figure 4-12.	STS-3 Transmitter Pulse Mask Corresponding to a Binary Zero	4-48
Figure 4-13.	STS-3 Transmitter Pulse Mask Corresponding to a Binary One	4-49
Figure 4-14.	STS-3 Eye Diagram Mask	4-50
Figure 5-1.	Example of Byte-Interleaving Sequence, STS-12	5-3
Figure 5-2.	Byte-Interleaving Example, Multiple Level Inputs.....	5-5
Figure 5-3.	Frame Synchronous Scrambler (Functional Diagram)	5-7
Figure 5-4.	Example of STS-1 and OC-N Signal Composition.....	5-9
Figure 5-5.	Switch Initiation Time Criteria.....	5-23
Figure 5-6.	Linear APS Switch – 1:n Architecture (in released position)	5-37
Figure 5-7.	Linear APS Switch – 1+1 Architecture (in released position)	5-37
Figure 5-8.	1:n Linear APS Architecture Example	5-45
Figure 5-9.	OC-N and OC-M Example.....	5-51
Figure 5-10.	External-Timing Mode Example	5-57
Figure 5-11.	Line-Timing Mode Example	5-59
Figure 5-12.	Loop-Timing Mode	5-59
Figure 5-13.	Through-Timing Mode Example.....	5-60
Figure 5-14.	Phase-Transient for Entry into Holdover.....	5-63
Figure 5-15.	OC-N Output Wander Time Deviation.....	5-65
Figure 5-16.	Time Deviation of Filtered Network Input to SONET NEs	5-66
Figure 5-17.	MTIE for SONET Clocks.....	5-68
Figure 5-18.	Time Deviation for SONET Clocks	5-69
Figure 5-19.	MTIE for Phase Transients from SONET Clocks	5-71
Figure 5-20.	DS1 Timing References Derived from a Single OC-N “Interface” Example	5-75
Figure 5-21.	DS1 Timing References Derived from Different OC-N “Interfaces” Example	5-75
Figure 5-22.	Time Deviation for Derived DS1	5-80

Figure 5-23.	Example Implementation of R5-211 and R5-212.....	5-96
Figure 5-24.	Example Implementation of R5-218 and R5-219.....	5-97
Figure 5-25.	Example Implementation of R5-220	5-97
Figure 5-26.	Category I DS1 and DS3 Jitter Transfer Mask.....	5-103
Figure 5-27.	Category II Jitter Transfer Mask	5-104
Figure 5-28.	SONET Category II Jitter Tolerance Mask	5-106
Figure 5-29.	Single Pointer Adjustment Test Sequence.....	5-111
Figure 5-30.	Maximum Rate Pointer Burst Test Sequence	5-112
Figure 5-31.	Phase Transient Pointer Burst Test Sequence	5-113
Figure 5-32.	Periodic VT1.5 Pointer Adjustment Test Sequence (26-1 Pattern)	5-116
Figure 5-33.	Periodic STS Pointer Adjustment Test Sequence (87-3 Pattern).....	5-117
Figure 5-34.	Periodic Pointer Adjustment Test Sequence (Continuous Pattern).....	5-118
Figure 5-35.	DS1 Mapping Phase Variation Limits.....	5-120
Figure 5-36.	DS3 Mapping Phase Variation Limits.....	5-121
Figure 5-37.	Single VT Pointer Adjustment Phase Variation Limits.....	5-123
Figure 5-38.	Single STS-1 Pointer Adjustment Phase Variation Limits	5-124
Figure 5-39.	Maximum Rate Pointer Burst Phase Variation Limits	5-125
Figure 5-40.	Phase Transient Pointer Burst Phase Variation Limits.....	5-126
Figure 5-41.	Periodic VT Pointer Adjustment Phase Variation Limits	5-127
Figure 5-42.	Periodic STS-1 Pointer Adjustment Phase Variation Limits	5-128
Figure 6-1.	General Defect Detection and Failure Declaration Model	6-15
Figure 6-2.	Maintenance Signals for SONET Layers	6-43
Figure 6-3.	STE Maintenance Signals.....	6-68
Figure 6-4.	LTE Maintenance Signals	6-69
Figure 6-5.	STS PTE (Asynchronous Mapping for DS3) Maintenance Signals	6-70
Figure 6-6.	STS PTE (VT-Structured STS-1 SPE) Maintenance Signals.....	6-71
Figure 6-7.	VT PTE (Asynchronous Mapping for DS _n into VT _x) Maintenance Signals, DS _n Interface	6-72
Figure 6-8.	VT PTE (Asynchronous Mapping for DS1 into VT1.5) Maintenance Signals, DS3 Interface with Embedded M13 Multiplex.....	6-73
Figure 6-9.	VT PTE (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals, DS1 Interface	6-74
Figure 6-10.	VT PTE (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals, DS1 Interface (Continued)	6-75
Figure 6-11.	VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces.....	6-76
Figure 6-12.	VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces (Continued).....	6-77
Figure 6-13.	DS0 PTE (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning	6-78
Figure 6-14.	SONET Maintenance Signals for Internal Equipment Failures	6-79

Figure 6-15.	Alarm Timing Requirements for Directly Detected Defects and Failures	6-80
Figure 6-16.	AIS Timing Requirements.....	6-81
Figure 6-17.	Derived RFI Timing Requirements	6-82
Figure 6-18.	RFI-V Signal Timing Requirements (for Byte-Synchronous DS1 Mapping)	6-82
Figure 6-19.	SONET PM Accumulation and Thresholding Model	6-95
Figure 6-20.	SONET PM Accumulation and Thresholding Model (Continued).....	6-96
Figure 6-21.	SONET PM Accumulation and Thresholding Model (Continued).....	6-97
Figure 6-22.	Intermediate-Path PM for Non-PJ-Related Parameters.....	6-129
Figure 6-23.	Examples of Intermediate-Path PM for STS PJ-Related Parameters ...	6-130
Figure 6-24.	Section Layer Diagnostic.....	6-139
Figure 6-25.	SONET Terminal Loopback.....	6-145
Figure 6-26.	SONET Facility Loopback	6-146
Figure 6-27.	DSn Terminal Loopback	6-148
Figure 6-28.	DSn Facility Loopback.....	6-148
Figure 8-1.	SONET Operations Communications: Example NE and Interface Types	8-1
Figure 8-2.	Example SONET Operations Communications Architecture	8-3
Figure 8-3.	Interactive Protocol Stacks for SONET Operations Communications.....	8-4
Figure 8-4.	Example Intra-site LAN and Point-to-Point DCC.....	8-5
Figure 8-5.	Example Operations Communications Network for a Survivable Ring...	8-5
Figure 8-6.	Operations Communications Functions.....	8-6
Figure 8-7.	SONET NSAP Format.....	8-19
Figure 8-8.	IDP Encoding	8-20
Figure 8-9.	TL1/X.25 – LAN Interworking	8-30
Figure 8-10.	TL1/X.25–LAN–DCC Interworking.....	8-31
Figure 8-11.	Example Routing Domain	8-35
Figure 8-12.	Craftsperson/NE Communications Network	8-38
Figure 8-13.	Manual TARP Adjacencies	8-52
Figure 8-14.	TARP Example.....	8-53

List of Tables

Table 2-1.	Reference Cable Lengths for Testing DS _n Interfaces	2-3
Table 3-1.	Line Rates for Standard SONET Interface Signals (N ≤ 48).....	3-2
Table 3-2.	STS Path Signal Label Assignments	3-35
Table 3-3.	STS Path Signal Label Assignments for Signals with Payload Defects.....	3-36
Table 3-4.	VT Signal Label Assignments	3-40
Table 3-5.	VT Size Indicator.....	3-78
Table 4-1.	Application Categories by Nominal Spectral Attributes	4-2
Table 4-2.	Attenuation and Group Delay Distortion as a Function of Frequency	4-9
Table 4-3.	LR OC-1 Optical Parameters	4-14
Table 4-4.	LR OC-3 Optical Parameters	4-15
Table 4-5.	LR OC-12 Optical Parameters	4-16
Table 4-6.	LR OC-48 Optical Parameters	4-17
Table 4-7.	IR OC-1 Optical Parameters	4-18
Table 4-8.	IR OC-3 Optical Parameters	4-19
Table 4-9.	IR OC-12 Optical Parameters	4-20
Table 4-10.	IR OC-48 Optical Parameters	4-21
Table 4-11.	SR OC-N Optical Parameters	4-22
Table 4-12.	STS-1 Electrical Interface Criteria.....	4-44
Table 4-13.	STS-3 Electrical Interface Criteria.....	4-47
Table 5-1.	Possible Starting Positions for an STS-Mc SPE in an OC-48 Signal	5-4
Table 5-2.	SONET Overhead Generating and Processing Criteria.....	5-12
Table 5-3.	Clearing Time Criteria for BER-based SF and SD Conditions	5-26
Table 5-4.	K1 Byte, Bits 1 through 4: Type of Request	5-30
Table 5-5.	Channel Number Code Assignments, K1 Bits 5 to 8 (and K2 Bits 1 to 4)	5-31
Table 5-6.	1:n Bidirectional Switching Example.....	5-46
Table 5-7.	Synchronization Status Message Definitions	5-53
Table 5-8.	Example of Reference Selection Using Synchronization Messages	5-91
Table 5-9.	Parameters for Network Interface Jitter Requirements	5-100
Table 5-10.	Category I Mapping Jitter Limits	5-108
Table 5-11.	Pointer Test Sequence Parameters.....	5-109
Table 5-12.	Jitter Due to Single Pointer Adjustments	5-111
Table 5-13.	Jitter Generation Limits for Periodic Pointer Adjustment Sequences	5-114
Table 6-1.	STS-1 Numbers in OC-N Signals (N = 3, 12, 48).....	6-3
Table 6-2.	STS Signal Label Mismatch Defect Conditions.....	6-34
Table 6-3.	VT Signal Label Mismatch Defect Conditions	6-37
Table 6-4.	RDI-P Bit Settings and Interpretation.....	6-56
Table 6-5.	RDI-V Bit Settings and Interpretation	6-60
Table 6-6.	Hierarchy of Near-end Failures	6-86
Table 6-7.	Hierarchy of Far-end Failures.....	6-87
Table 6-8.	PM Register Sizes and Default Thresholds	6-102

Table 6-9.	Section BIP Errors to Trigger a Section SES	6-108
Table 6-10.	Line BIP Errors to Trigger a Line SES (Near-End and Far-End)	6-110
Table 6-11.	STS Path BIP Errors to Trigger an STS Path SES (Near-End and Far-End)	6-113
Table 6-12.	VT Path BIP Errors to Trigger a VT Path SES (Near-End and Far-End)	6-118
Table 6-13.	Section Layer PM Accumulation During Defects	6-124
Table 6-14.	Near-End Line Layer PM Accumulation During Defects	6-125
Table 6-15.	Near-End STS Path Layer PM Accumulation During Defects	6-126
Table 6-16.	Near-End VT Path Layer PM Accumulation During Defects	6-127
Table 6-17.	Far-End Line Layer PM Accumulation During Defects	6-127
Table 6-18.	Far-End STS Path Layer PM Accumulation During Defects	6-128
Table 6-19.	Far-End VT Path Layer PM Accumulation During Defects	6-128
Table 8-1.	TARP PDU Fields	8-41
Table 8-2.	TARP Types	8-42
Table 8-3.	TARP Timers	8-44
Table 8-4.	Provisionable TARP PDU Fields	8-50

Preface

This Preface contains important information about Bellcore's GR process in general, as well as important information about this document.

Bellcore's GR Process

Generic Requirements documents (GRs) provide Bellcore's view of proposed generic criteria for telecommunications equipment, systems, or services, and involve a wide variety of factors, including interoperability, network integrity, funding participant expressed needs, and other input.

Bellcore's GR process implements Telecommunications Act of 1996 directives relative to the development of industry-wide generic requirements relating to telecommunications equipment, including integral software and customer premises equipment. Pursuant to that Act, Bellcore invites members of the industry to fund and participate in the development process for such GRs. Invitations to fund and participate are issued monthly in the *Bellcore Digest of Technical Information*, and posted on Bellcore's web site at <http://www.bellcore.com/DIGEST>.

At the conclusion of the GR development process, Bellcore publishes the GR, which is available by subscription. The subscription price entitles the purchaser to receive that issue of the GR (GR-CORE) along with any Issues List Report (GR-ILR) and Revisions, if any are released under that GR project. ILRs contain any technical issues that arise during GR development that Bellcore and the funding participants would like further industry interaction on. The ILR may present issues for discussion, with or without proposed resolutions, and may describe proposed resolutions that lead to changes to the GR. Significant changes or additional material may be released as a Revision to the GR-CORE.

Bellcore may also solicit general industry nonproprietary input regarding such GR material at the time of its publication, or through a special Industry Interaction Notice appearing in the *Bellcore Digest of Technical Information*. While unsolicited comments are welcome, any subsequent work by Bellcore regarding such comments will depend on funding support for such GR work. Bellcore will acknowledge receipt of comments and will provide a status to the submitting company.

About GR-253-CORE

A. Funders of GR-253-CORE Issue 2, Revision 2

Alcatel Network Systems
Ameritech
Bell Atlantic
BellSouth
Fujitsu Network Communications, Inc.
SBC Communications, Inc.
U S WEST

B. Relative Maturity Level

This is a mature technology and requirements reflect maintenance mode. Throughout the current GR-253-CORE document (i.e., GR-253-CORE, Issue 2 plus Revisions 1 and 2), the criteria are considered stable except as indicated in the text. Issues affecting these criteria will appear in GR-253-ILR, Issue 2C, scheduled for release in February 1999.

C. GR-253-CORE Plans

This document is expected to be reissued in 2000.

To Submit Comments

When submitting comments, please include the GR document number, and cite any pertinent section and requirement number. In responding to an ILR, please identify the pertinent Issue ID number. Please provide the name and address of the contact person in your company for further discussion.

Comments should be submitted by **August 1, 1999**.

Send comments to:

Bellcore — GR-253-CORE
Anna Reidy
SONET Project Management
331 Newman Springs Road
Red Bank, NJ 07701-5699

Phone: (732) 758-2817
FAX: (732) 758-4177
E-Mail: areidy@notes.cc.bellcore.com

1. Introduction

GR-253-CORE, Issue 2, along with GR-253-CORE, Issue 2, Revision 1 and GR-253-CORE, Issue 2, Revision 2 contain Bellcore's view of Synchronous Optical Network (SONET) generic criteria. In general, the criteria in GR-253-CORE, Issue 2 expanded on, and in some cases changed the criteria contained in GR-253-CORE, *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria*, Issue 1, December 1994. In turn, Revisions 1 and 2 to GR-253-CORE, Issue 2, expand on, and in some cases changes the criteria in several sections of GR-253-CORE, Issue 2.

GR-1377-CORE, *SONET OC-192 Transport System Generic Criteria*, provides proposed generic criteria in support of SONET OC-192. There presently are no plans to merge that document with GR-253-CORE.

The criteria contained herein are intended to advise the telecommunications industry of Bellcore's view of proposed generic requirements. The criteria reflect recent changes in requirements advocated by the various committees and subcommittees of standards organizations such as the International Telecommunication Union - Telecommunication Standardization Sector (ITU-T), ANSI-accredited Committee T1, and the Electronic Industries Association/Telecommunications Industries Association (EIA/TIA) and recent agreements reached in the SONET Interoperability Forum (SIF).

1.1 Requirements Terminology

The following requirements terminology is used in GRs:

- Requirement — Feature or function that, in Bellcore's view, is *necessary* to satisfy the needs of a typical service provider. Failure to meet a requirement may cause application restrictions, result in improper functioning of the product, or hinder operations. A Requirement contains the words *shall* or *must* and is flagged by the letter "R."
- Conditional Requirement — Feature or function that, in Bellcore's view, is *necessary in specific applications*. If a service provider identifies a Conditional Requirement as necessary, it shall be treated as a requirement for the application(s). Conditions that may cause the Conditional Requirement to apply include, but are not limited to, certain service provider application environments, elements, or other requirements, etc. A Conditional Requirement is flagged by the letters "CR."
- Objective — Feature or function that, in Bellcore's view, is *desirable* and may be required by a service provider. An Objective represents a goal to be achieved. An Objective may be reclassified as a Requirement at a specified date. An objective is flagged by the letter "O" and includes the words *should*, *it is desirable*, or *it is an objective*.

- Conditional Objective — Feature or function that, in Bellcore's view, is *desirable in specific applications* and may be required by a service provider. It represents a goal to be achieved in the specified Condition(s). If a service provider identifies a Conditional Objective as necessary, it shall be treated as a requirement for the application(s). A Conditional Objective is flagged by the letters "CO."
- Condition — The circumstances that, in Bellcore's view, will cause a Conditional Requirement or Conditional Objective to apply. A Condition is flagged by the letters "Cn."

1.2 Requirement Labeling Conventions

As part of Bellcore's new GR Process, proposed requirements and objectives are labeled using conventions that are explained in the following two sections.

1.2.1 Numbering of Requirement and Related Objects

Each Requirement, Objective, Condition, Conditional Requirement, and Conditional Objective object is identified by both a local and an absolute number. The local number consists of the object's document section number and its sequence number in the section (e.g., **R3-1** is the first Requirement in Section 3). The local number appears in the margin to the left of the Requirement. A Requirement object's local number may change in subsequent issues of a document if other Requirements are added to the section or deleted.

The absolute number is a permanently assigned number that will remain for the life of the Requirement; it will not change with new issues of the document. The absolute number is presented in brackets (e.g., [2]) at the beginning of the requirement text.

Neither the local nor the absolute number of a Conditional Requirement or Conditional Objective depends on the number of the related Condition(s). If there is any ambiguity about which Conditions apply, the specific Condition(s) will be referred to by number in the text of the Conditional Requirement or Conditional Objective.

References to Requirements, Objectives, or Conditions published in other Generic Requirements documents will include both the document number and the Requirement object's absolute number. For example, **R2345-12** refers to Requirement [12] in GR-2345.

1.2.2 Requirement, Conditional Requirement, and Objective Object Identification

A Requirement object may have numerous elements (paragraphs, lists, tables, equations, etc.). To aid the reader in identifying each part of the requirement, an ellipsis character (...) appears in the margin to the left of all elements of the Requirement.

This GR helps establish a foundation for interoperability between different implementations of the functions described herein. It is important to note, however, that a number of optional features are included, both in this document and the standards documents that are referenced. Criteria in this GR relating such features are denoted as conditional requirements. In selecting a consistent set of optional features, the following situations have to be taken into account:

- The domain of a single network provider when using a single supplier
- The domain of a single network provider when using multiple suppliers
- Interoperability between domains of different network providers.

1.3 Revision History

This section gives a high-level view of the changes in GR-253-CORE, Issue 2 with respect to GR-253-CORE, Issue 1, and of the subsequent changes that were made in GR-253-CORE, Issue 2, Revisions 1 and 2. Many of these changes are related to issues that were discussed in GR-253-ILR, Issues 1A, 2A and 2B, or are based on industry comments and agreements reached in standards bodies or the SIF after those documents were issued. In addition, changes were made in several sections to further clarify the criteria, and to eliminate redundant criteria. In general, the changes that were made in Issue 2 (with respect to Issue 1) were not marked with change bars.¹ On the other hand, all of the substantive changes that were made in Issue 2, Revision 1 (with respect to Issue 2) were marked with change bars in the outside margins of the document, as were all of the substantive changes in Issue 2, Revision 2 (with respect to Issue 2 and Issue 2, Revision 1). On the other hand, all of the major sections that were revised in Revision 1 were also revised in Revision 2, and only the change bars associated with the changes that first appeared in Revision 2 currently appear. That is, none of the change bars that appeared in Revision 1 have been retained in Revision 2, and therefore they do not appear in the complete document.

The following lists describe the major changes by section. Note that criteria with absolute numbers from [903] to [1012] are criteria that first appeared in Issue 2 of this document, while criteria with absolute numbers from [1013] to [1043] first appeared in Issue 2, Revision 1, and those with absolute numbers from [1044] to [1099] first appeared in Issue 2, Revision 2. Also, a version number (e.g., [96v2]) has been added to the absolute numbers for criteria with substantive changes (including cases where a previously existing requirement was split into two or more criteria). Finally, the absolute numbers assigned to criteria that appeared in a previous issue, but that have now been removed, are not reused.

1. Note that some older copies of GR-253-CORE, Issue 2 may also contain change bars; however, those change bars should not be considered complete and should be ignored.

1.3.1 Changes Between GR-253-CORE Issues 1 and 2

As itemized below, GR-253-CORE, Issue 2 contains at least some changes (from GR-253-CORE, Issue 1) in most of its major sections.

- Section 2, *Network Compatibility*, changes include:
 - Clarification of the reference cable listings in Section 2.1.1.1
 - Section 3, *Rates and Formats*, changes include:
 - Clarification, in Section 3.2, that an NE can set certain unused bits and bytes to non-zero values and still conform to the objective in that section
 - Modification of the criteria in Sections 3.3.2.3 and 3.3.3 on sending an all-zeros STS or VT SPE to allow an STS or VT PTE that has not yet been assigned a signal on its low-speed side to generate a path trace message and possibly a non-zero signal label on its high-speed side
 - Elimination of the PDI-V indication which used a code in the VT Path Signal Label to indicate the existence of a defect in the embedded payload (ILR Issue ID 253-5)
 - Change the bits reserved for RDI-P from bits 5 and 6 of the G1 byte, to bits 5, 6 and 7 of G1 (ILR Issue ID 253-25, see Sections 3.3.2.3 and 6.2.1.3.2)
 - Change the bits reserved for RDI-V from bits 4 and 8 of the V5 byte, to bits 5, 6 and 7 of the Z7 byte and bit 8 of V5 (ILR Issue ID 253-25, see Sections 3.3.3 and 6.2.1.3.3)
 - Change the bit reserved for RFI-V from bit 8 of the Z7 byte, to bit 4 of V5 (ILR Issue ID 253-25, see Sections 3.3.3 and 6.2.1.3.3)
 - The addition of criteria in Sections 3.4.1.2, 3.4.1.3, 3.4.1.4, and 3.4.2.1 on DS1, DS1C, DS2, and DS3 bit stuffing jitter transfer (ILR Issue ID 253-11)
 - The addition of a requirement in Section 3.4.2.1 on DS3 bit stuffing jitter generation (ILR Issue ID 253-6)
 - The addition of information and criteria in Sections 3.5.1.5 and 6.2.1.2.2 related to the possible responses of non-STs PTE (e.g., a pointer processor in LTE) to an incoming all-ones STS pointer word
 - The addition of information and criteria in Section 3.5.2.5 and 6.2.1.2.3 related to the possible responses of non-VT PTE (e.g., a VT pointer processor in STS PTE) to an incoming all-ones VT pointer word
 - Section 4, *Physical Layer*, changes include:
 - The addition of a reference in Section 4.1 to ANSI standards for applications not covered in this GR (e.g., SR-0 applications, OC-24)
-

- Clarification in Section 4.2.5 regarding the possible need for attenuators in Long Reach OC-48 applications
- The addition of information in Section 4.4 regarding the expected pulse amplitudes for STS-1 electrical signals meeting the wideband power level requirement
- Section 5, *Network Element Architectural Features*, changes include:
 - The addition of further restrictions on the possible starting positions of an STS-Mc SPE in an OC-N signal (ILR Issue ID 253-2, see Section 5.1.1)
 - The addition of a reference to GR-499-CORE for criteria related to circuit pack protection switching performance (see Section 5.3)
 - The addition of criteria in Section 5.3.3.3 to encourage error-free manually-initiated switches (when possible)
 - The addition of criteria in Section 5.3.4 to reduce the chance of protection switch oscillations in certain double failure scenarios (ILR Issue ID 253-8)
 - Clarification of the bit assignments for the K2 byte in Section 5.3.5.2.1
 - The addition of information and criteria in Section 5.3.5.5 regarding the acceptance of both "high" and "low priority" SD and SF codes on the incoming K1 byte as valid requests
 - The addition of information in Section 5.3.5.5 regarding the expected operation of a linear APS system when one NE is provisioned to use nonrevertive switching, and the other is provisioned to use revertive switching
 - Clarifications and modifications throughout Section 5.4 on synchronization in response to developments in ANSI and the issuance of GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria*
 - The addition of information in Section 5.4.3.4 discouraging the use of through-timing at NEs that terminate the SONET Line layer
 - The addition of a new section, Section 5.6.2.3.7, with criteria related to DS_n pointer adjustment jitter generation at an NE where the DS_n does not appear at an external DS_n interface (ILR Issue ID 253-12)
 - Clarification of the jitter tolerance criteria in Section 5.6.2.2.2 for OC-48 interfaces
- Section 6, *SONET Network Element Operations Criteria*, changes include:
 - The addition of criteria in Section 6.1.2 on the numbering of STS-1 and STS-Mc modules in an OC-N signal (ILR Issue ID 253-4)
 - Clarification of the LOS detection criteria in Section 6.2.1.1.1
 - Clarification and the revision of the DS_n OOF detection criteria in Section 6.2.1.1.2

- Clarification in Section 6.2.1.1.3 that LOP defects are not detected when equipment that does not terminate a path (e.g., LTE that processes the STS pointers) is relaying an incoming all-ones pointer
- Clarification that an NE using 1+1 unidirectional linear APS may (but is not required to) monitor the APS channel for APS-related defects and failures described in Section 6.2.1.1.6
- Correction of the text in Section 6.2.1.1.6.D related to the actions performed by an NE when it detects a Far End Protection Line defect
- Clarification of the tables on the detection of UNEQ and PLM defects (see Tables 6-2 and 6-3)
- The addition of requirements in Section 6.2.1.2.1 concerning STE's generation of AIS-L when it detects the failure of LTE supporting provisioned line origination functions
- Clarification and the revision of the DS_n AIS generation and detection criteria in Section 6.2.1.2.4
- Revision of the criteria in Section 6.2.1.3.1 related to the removal of RDI-L by an NE using linear APS, and to the detection and termination times for RDI-L defects
- Revision of the minimum assertion times for RDI-L, RDI-P, and RDI-V signals (ILR Issue ID 253-32, see Sections 6.2.1.3.1, 6.2.1.3.2 and 6.2.1.3.3)
- The addition of information in Sections 6.2.1.3.2 and 6.2.1.3.3 regarding a number of open issues related to the triggers for RDI-P and RDI-V generation, and the modification of the criteria to support both "one-bit" and "enhanced" versions of those signals (until the open issues are resolved)
- Clarification of the RFI-V signal detection criteria in Section 6.2.1.3.3
- Clarification and revision of the DS_n RAI and RDI generation and detection criteria in Section 6.2.1.3.4
- Clarification and separation of the criteria in Section 6.2.1.4.1 on generating and detecting PDI-P signals
- Clarification of the trunk conditioning criteria in Section 6.2.1.6
- Simplification of Figures 6-3 through 6-13 (and the elimination of GR-253-CORE, Issue 1 Figure 6-14) on maintenance signal generation
- The addition of a new figure, Figure 6-14, illustrating the generation of AIS-L, AIS-P, and AIS-V upon the detection of an internal equipment failure
- Clarification and revision (for consistency) of the criteria in Section 6.2.1.8 on declaring failures in the presence of other failures

- Separation of Section 6.2.1.8 into seven subsections on various topics, including Section 6.2.1.8.7 on the declaration and clearing of failures when a signal is being non-intrusively monitored
- Clarification of a number of the PM parameter definitions in Section 6.2.2
- Clarification of the general PM accumulation and thresholding criteria in Section 6.2.2.1
- The addition of information and criteria in Section 6.2.2.1 regarding the accumulation of PM parameter at the end of a period when the possibility exist for entry into or exit from unavailable time during the first 10 seconds of the next period
- The addition of information in Section 6.2.2.1 regarding the minimum PM register sizes for the various CV parameters
- Clarification of the Section-layer PM accumulation criteria in Section 6.2.2.3.2 for the case of an incoming drop-side signal
- Clarification of the criteria in Section 6.2.2.8 on performance monitoring during troubles, and simplification of the table (Table 6-12 in Issue 1) on the same topic (by dividing it into separate tables for each layer, see Tables 6-13 through 6-19)
- Clarification of the intermediate path PM definitions and criteria in Section 6.2.2.9
- Clarification and revision of the STS Path Trace criteria in Section 6.2.3.2.3.A, including the addition of criteria supporting the use of an "expected" incoming path trace
- The addition of criteria regarding STS Path Signal Label diagnostics to support the use of PDI-P (see Section 6.2.3.2.3.B)
- Reclassification of the function described in Section 6.2.3.3.1 of Issue 1 from a "SONET terminal loopback" to a "diagnostic" (see Section 6.2.3.2.2), and the addition of a new SONET terminal loopback description
- Clarification of the SONET facility, DS_n terminal, and DS_n facility loopback criteria (including new figures) in Section 6.2.3.3
- Section 7, *Other Generic Criteria*, changes include:
 - Removal of high temperature safety label requirements in Section 7.4.1, and the addition of a reference to GR-499-CORE, *Transport Systems Generic Requirements (TSGR): Common Requirements*, for criteria on that topic (ILR Issue ID 253-13)
- Section 8, *SONET Operations Communications*, changes include:
 - Modification of the criteria in Section 8.3.1.2 concerning the physical interface for LANs

- Clarification of TARP PDU processing in Section 8.7.5, and the loop detection and propagation procedures in Sections 8.7.5.7 and 8.7.5.8
- Modification of the TARP pseudocode in Section 8.7.10
- Appendix A, *Requirement-Object List*, has been changed to reflect the new and revised criteria
- Appendix B, *Fiber Optic Transmission System Design Worksheets*, has no changes
- Appendix C, *SONET Operations Communications Lower Layers Protocol Profile*, has no major changes
- Appendix D, *SONET Operations Communications Upper Layers Protocol Profile*, has no major changes
- The *References* section was updated
- A number of definitions, including those for STE, LTE, STS PTE, and VT PTE were revised or corrected in the *Glossary* section
- The *Requirement-Objects Index* section was updated.

1.3.2 Changes Between GR-253-CORE Issue 2 and Revision 1 to Issue 2

As itemized below, GR-253-CORE, Issue 2, Revision 1 contains technical changes (from GR-253-CORE, Issue 2) in Sections 5, 6 and 8, and Appendix C. In addition, the Requirement-Object List (Appendix A), References section, and Requirement-Objects Index have been updated to reflect those changes.

- Section 5, *Network Element Architectural Features*, changes include:
 - The addition of information related to the possible sizes for an STS-Mc SPE in an OC-N signal (ILR Issue ID 253-38, see Section 5.1.1)
 - The addition of information and criteria related to the acceptance of off-frequency SONET signals (ILR Issue ID 253-45, see Section 5.2.1)
 - Clarification of the intended definition of the word "pass" in the text and criteria related to orderwire (ILR Issue ID 253-46, see Section 5.2.2.2)
 - The addition of information related to the transport of the extra traffic channel in a system supporting 1:n linear APS (ILR Issue ID 253-9, see Section 5.3.5.5)
 - The addition of information related to the possible reinitiation of preempted external linear APS switch requests (ILR Issue ID 253-47, see Section 5.3.6.1)
 - Clarification of the use of the word "interface" in the various synchronization-related criteria (see Section 5.4)

- Clarification of the definition of a SONET Minimum Clock (SMC), and the revision of a number of synchronization-related criteria to reflect that definition (ILR Issue ID 253-56, see Sections 5.4.1, 5.4.4.2, 5.4.4.2.2 and 5.4.4.3.3)
- The acknowledgment of ongoing standards work related to synchronization status messages (see Section 5.4.2)
- Revision of the text related to the possibility of provisioning the working and protection lines in a system with line APS as separate timing references (see Section 5.4.3.2)
- Clarification of the intent of the holdover criteria in Section 5.4.4.2.2
- Clarification of the intent of the pull-in/hold-in "settling time" criteria in Section 5.4.4.2.3
- Clarification of the applicability of the SMC and stratum 3 clock wander transfer criteria (ILR Issue ID 253-55, see Section 5.4.4.2.4)
- The addition of information related to the testing of a SONET NE's clock against the wander generation criteria in Section 5.4.4.3.2
- Revision of the phase transient criteria in Section 5.4.4.3.3 to be more consistent with the corresponding criteria in GR-1244-CORE, and to specifically include timing-related working/protection line switches in the list of synchronization rearrangement operations that are allowed to cause such transients (ILR Issue ID 253-58)
- The separation of the criteria and text related to jitter and errors during synchronization rearrangement operations (see Section 5.4.4.3.4) from the phase transient criteria (Section 5.4.4.3.3)
- The revision of the requirement related to the maximum rate of frequency change during holdover recovery (ILR Issue ID 253-49, see Section 5.4.4.3.5)
- The addition of criteria related to an NE's tolerance of jitter on its incoming timing reference signals (ILR Issue ID 253-57, see Section 5.4.4.3.6)
- Clarification and addition of text and criteria related to an NE's ability to derive DS1 signals from one or more incoming OC-N signals for timing distribution purposes (ILR Issue IDs 253-50 and 253-51, see Section 5.4.5.1)
- The addition of information related to the input signals for use in derived DS1 wander generation and jitter generation tests (ILR Issue IDs 253-52 and 253-53, see Section 5.4.5.1)
- The addition of information related to derived DS1 source switching and timing reference switching based on synchronization status messages (see Sections 5.4.5.2.1 and 5.4.6.4)

- The addition of information related to the manual switch command currently defined to support switching between provisioned timing references (see Section 5.4.6)
 - The addition of criteria for considering an OC-N timing reference as failed (ILR Issue ID 253-59, see Section 5.4.6.1)
 - The addition of information related to revertive and nonrevertive switching between timing references (see Section 5.4.6.3)
 - The addition of text and criteria related to holdoff times for certain actions that are required to be triggered by changes in an NE's received synchronization status messages (ILR Issue IDs 253-60 and 253-62, see Sections 5.4.6.4 and 5.4.7.3.1)
 - Clarification of the objective on retrieving synchronization status messages (ILR Issue ID 253-61, see Section 5.4.7.1)
 - The addition of information related to measurement methods for systems with non-linear jitter transfer characteristics (ILR Issue ID 253-64, see Section 5.6.2.1)
 - The addition of text and criteria related to the generation of pointer adjustments in the presence of a constant frequency offset (ILR Issue ID 253-15, see Section 5.6.2.3.2)
 - The addition of text and criteria regarding phase variations for DS3 payload signals transported on SONET networks (ILR Issue ID 253-67, see Section 5.7).
 - Section 6, *SONET Network Element Operations Criteria*, changes include:
 - Modification of the criteria regarding the effect of incoming all-ones signal labels on existing PLM or UNEQ defects (ILR Issue ID 253-71, see Sections 6.2.1.1.8.B and 6.2.1.1.8.D)
 - Modification of the RDI generation and detection criteria to align with ANSI T1.231-1997 (ILR Issue IDs 253-23, 253-34, 253-74, 253-75, 253-76 and 253-77, see Sections 6.2.1.3.2 and 6.2.1.3.3)
 - The addition of DS_n AIS to the hierarchy of near-end failures table in Section 6.2.1.8.2 (ILR Issue ID 253-31)
 - Modification of the criteria related to the support of invalid data flags for various SONET PM parameter registers (ILR Issue ID 253-78, see Section 6.2.2.1)
 - Modification of the minimum threshold register size criteria for the SEFS, SES and UAS parameters defined at the various SONET layers (see Section 6.2.2.1)
 - Clarification of the intent and applicability of the objectives to support SONET Physical layer PM (see Section 6.2.2.2.2)
-

- Revision of the values of K in the SONET Section, Line and STS Path SES parameter definitions to align with ANSI T1.231-1997 (see Sections 6.2.2.3.1, 6.2.2.4.1 and 6.2.2.5.1)
- The addition of text and criteria to support the accumulation of SONET Section layer PM parameters at drop-side interfaces (ILR Issue ID 253-85, see Section 6.2.2.3.2)
- The addition of references to the pointer justification-related PM parameter definitions in ANSI T1.231-1997 (see Sections 6.2.2.4.1 and 6.2.2.5.1)
- Revision of the STS and VT Path layer ES, SES and FC parameter definitions and accumulation criteria to align with ANSI T1.231-1997 (ILR Issue IDs 253-74, 253-87, 253-88 and 253-89, see Sections 6.2.2.5.1, 6.2.2.5.2, 6.2.2.6.1, 6.2.2.6.2 and 6.2.2.8)
- Revision of the far-end STS and VT Path layer PM criteria to indicate that those parameters are not accumulated when an UNEQ defect is present at the same layer (ILR Issue ID 253-74, see Section 6.2.2.8)
- Clarification and revision of the intermediate-path PM criteria to align with ANSI T1.231-1997 (ILR Issue ID 253-90, see Section 6.2.2.9)
- Revision of the description of the 16-byte path trace format used in SDH (see Section 6.2.3.2.3.A)
- Clarification of the effect of an incoming DS_n OOF (see Figures 6-5 and 6-7).
- Section 8, *SONET Operations Communications*, changes include:
 - Various changes that were based on SIF-GN-9608-06R5, *SIF Implementation Agreements for SONET Operations Communications*, including:
 - The addition of an objective to permit LAN-based OS-NE communications (ILR Issue ID 253-96, see Sections 8.2.1, 8.4, 8.4.4 and 8.5.3)
 - The addition of a requirement which specifies the upper limit of LT1 messages size to 4096 bytes for TL1 over any protocol stack or transport mechanism, (ILR Issue ID 253-97, see Section 8.3.7.5)
 - The addition of a requirement that case shall be ignored for all TARP-related TID/NET mappings (ILR Issue ID 253-97, see Section 8.7)
 - The addition of text clarifying the conditions under which a TARP PDU is considered invalid and may be discarded (ILR Issue ID 253-97, see Sections 8.7.2 and 8.7.2.4)
 - The addition of a requirement stating that the URC bit shall be ignored upon receipt of TARP PDUs (ILR Issue ID 253-97, see Section 8.7.2.4)

- The removal of support for Level 2 TDC (ILR Issue ID 253-97, see Sections 8.7.3 and 8.7.5.6.3)
 - The revision of a requirement to reserve zero in the tar-seq field to indicate that a reset has occurred (ILR Issue ID 253-97, see Section 8.7.5)
 - The addition of a requirement which makes the tar-tor field in a TARP Type 1 or Type PDU optional (ILR Issue ID 253-97, see Section 8.7.5.1)
 - The revision of the TARP pseudocode to reflect the changes listed above (ILR Issue ID 253-97, see Section 8.7.10)
 - The addition of an objective to support X.500-based directory services for the name/address translation service (see Sections 8.3 and 8.3.7.4)
 - The removal of a note indicating that a forthcoming revision to GR-828-CORE would align it with GR-253-CORE (ILR Issue ID 253-95, see Sections 8.3.3.1 and 8.5.2)
 - The revision of the description of the Craftsperson/NE interface, and the addition of two objectives based on SIF-009-1997 (see Sections 8.6 and 8.6.2).
 - Appendix A, *Requirement-Object List*, was changed to reflect the new and revised criteria.
 - Appendix C, *SONET Operations Communications Lower Layers Protocol Profile* changes include:
 - The addition of a note stating that the tar-tor field is optional (see Section C.8.2)
 - The addition of a conformance statement on ignoring the URC bit upon receipt of TARP PDUs (see Section C.8.3.2)
 - The addition of a conformance statement to support the LDB Entry Timer (ILR Issue ID 253-98, see Section C.8.4)
 - The addition of conformance statements to support manual provisioning of the LDB Entry Timer and the LDB Flush Timer (ILR Issue ID 253-98, see Section C.8.5)
 - The addition of LDB Entry Timer Parameters conformance statements (ILR Issue ID 253-98, see Section C.8.5.1)
 - The addition of LDB Flush Timer Parameters conformance statements (ILR Issue ID 253-98, see Section C.8.5.2)
 - The addition of a conformance statement that tar-seq is a provisionable TARP PDU field (ILR Issue ID 253-98, see Section C.8.5.3).
 - The *References* section was updated.
 - The *Requirement-Objects Index* section was updated.
-

1.3.3 Changes In Revision 2 to GR-253-CORE Issue 2

As itemized below, Revision 2 to GR-253-CORE, Issue 2, contains technical changes (from GR-253-CORE, Issue 2 and GR-253-CORE, Issue 2, Revision 1) in Sections 3, 5, 6 and 8, and Appendix C. In addition, the Requirement-Object List (Appendix A), References section, and Requirement-Objects Index have been updated to reflect those changes.

- Section 3, *Rates and Formats*, changes included:
 - The addition of information related to the possible sizes for an STS-Mc SPE in an OC-N signal (ILR Issue ID 253-38, see Section 3.2.3)
 - The addition of the "HDLC-over-SONET Mapping" and "O.181 Test Signal (TSS1 to TSS3) Mapping" STS path signal label codes to Table 3-2 (ILR Issue ID 253-39)
 - The addition of text and criteria defining an HDLC-over-SONET mapping for use in the transport of HDLC-framed traffic in STS SPEs (ILR Issue ID 253-39, see Sections 3.4.2.3, 3.4.3.1.5 and 3.4.3.2.2)
 - The addition of information regarding the ability (or inability) of an NE that is performing all-ones pointer relay to also relay the associated all-ones STS or VT SPE [see Sections 3.5.1.5 and 3.5.2.5 (and Sections 6.2.1.2.2 and 6.2.1.2.3)].
- Section 5, *Network Element Architectural Features*, changes included:
 - The addition of information related to the effect of the APS architecture supported by an NE on its processing of various overhead bits and bytes (see Section 5.2.1)
 - The addition of information related to the limitations of the required orderwire channel protection scheme in configurations where two NEs supporting linear APS are connected via diversely-routed working and protection lines that include one or more regenerators (ILR Issue ID 253-99, see Section 5.2.2.2)
 - The revision of the detection and clearing criteria for linear APS SD and SF conditions, to be consistent with the corresponding ITU-T specifications (ILR Issue IDs 253-101 and 253-102, see Sections 5.3.3.2 and 5.3.4)
 - The revision of an existing requirement to indicate that the WTR period must be user-provisionable on a per-protection line (or per-protection group) basis (ILR Issue ID 253-103, see Section 5.3.4)
 - The revision of the required functionality for the "Lockout a Working Channel" linear APS control command, to be consistent with the corresponding ITU-T specifications (ILR Issue ID 253-104, see Section 5.3.6.2)
 - The addition of new synchronization status message codes to Table 5-7 for signals whose timing is "Transit Node Clock Traceable" or "Stratum 3E Traceable" (see Section 5.3.6.2)

- The revision and addition of several criteria to indicate that NEs are always required to generate appropriate synchronization status messages on their outgoing SONET signals, and that they need to provide at least the "capability to process" those messages (e.g., as a user-provisionable option) on certain of their incoming signals (ILR Issue ID 253-107, see Section 5.4.2)
- The revision and addition of text and several criteria to indicate that an NE that supports line APS is allowed to support the provisioning of working line 1 and the protection line at a single SONET "interface" as separate references (ILR Issue ID 253-105, see Sections 5.4.3.2 and 5.4.6.1)
- The revision of several criteria to clarify that the input and output wander transfer TDEV masks were intended to allow up to 0.2 dB of gain in the pass-band of the clock (ILR Issue ID 253-108, see Section 5.4.4.2.4)
- The revision of the description of the input signal used for testing the conformance of a line-timed or through-timed NE to the wander generation requirement in Section 5.4.4.3.2 (ILR Issue ID 253-52)
- The addition of text and criteria to support two new synchronization-related switching commands (ILR Issue ID 253-113, see Sections 5.4.5.1 and 5.4.6)
- The revision of the description of the input signal used for testing the conformance of a SONET NE to the derived DS1 jitter generation requirement in Section 5.4.5.1 (ILR Issue ID 253-52)
- The replacement of the derived DS1 wander generation TDEV mask in Figure 5-22 (ILR Issue ID 253-109)
- The clarification of the criteria related to the Manual Reference Switch command (ILR Issue IDs 253-110 and 253-112, see Section 5.4.6)
- The clarification of the objective related to on-demand reporting of an NE's current synchronization status messages (ILR Issue ID 253-107, see Section 5.4.7.1)
- The addition of criteria to allow a user to disable the automatic generation of the DUS synchronization status message by an externally timed NE (ILR Issue ID 253-115, see Section 5.4.7.3.1)
- The revision of the DS1 and DS3 periodic pointer adjustment phase variation criteria, to align with the corresponding ANSI specifications (see Section 5.7.2.3)
- Section 6, *SONET Network Element Operations Criteria*, changes included:
 - Modification of the criteria regarding the detection of LOP defects and the termination of AIS defects (ILR Issue ID 253-116, see Sections 6.2.1.1.3, 6.2.1.2.2 and 6.2.1.2.3)
 - The addition of criteria related to the alarming of BER-based SF and SD conditions (ILR Issue ID 253-117, see Sections 6.2.1.1.6.E)

- The removal of a number of notes that previously indicated that the insertion of AIS downstream upon detection of an UNEQ or PLM defect was under study (ILR Issue ID 253-118)
- The correction of note "c" in Tables 6-2 and 6-3 regarding the effect of incoming all-ones signal labels on existing PLM or UNEQ defects
- The addition of a new section defining an STS Path Trace Identifier Mismatch (TIM-P) defect, and the modification of criteria and text throughout Section 6 to reflect the existence of that defect (ILR Issue ID 253-23, see Section 6.2.1.1.9)
- The addition of information related to a proposed "Automatic In-Service Provisioning" (AISP) state, and the corresponding modification of several criteria on the generation of autonomous messages to an OS (ILR Issue ID 253-129, see Sections 6.2.1.8, 6.2.1.8.2 and 6.2.2.1)
- The modification of text and criteria to clearly differentiate between the two definitions of the word "failure" that were previously used in Section 6.2.1.8 (ILR Issue ID 253-122)
- The revision of an objective so that it allows for the autonomous reporting of failures declared for non-intrusively monitored signals (ILR Issue ID 253-123, see Section 6.2.1.8.7)
- The revision of text and a number of criteria related to the accumulation of SONET Physical layer PM parameters, to align with the corresponding ANSI specifications (ILR Issue IDs 253-82, 253-83, 253-84 and 253-124, see Sections 6.2.2.1 and 6.2.2.2)
- The removal of text and criteria related to a previously allowed option for processing PM parameters near time period boundaries (ILR Issue ID 253-81, see Section 6.2.2.1)
- The addition of values to Table 6-8 for the minimum CV register and CV threshold register sizes (ILR Issue ID 253-80)
- The replacement of text and criteria related to the previously defined STS and VT PJ PM parameters with new text and criteria that are aligned with the corresponding ANSI specifications (ILR Issue ID 253-28, see Sections 6.2.2.4, 6.2.2.5, 6.2.2.6 and 6.2.2.9)
- The addition of text and criteria to indicate that the supplier must clearly document the number or percentage of the paths for which an NE can simultaneously perform intermediate-path PM (ILR Issue ID 253-126, see Section 6.2.2.9)
- The modification of the SONET Physical layer diagnostics criteria to be consistent with the revised Physical layer PM criteria (ILR Issue ID 253-82, see Section 6.2.3.2.1)

- The revision of text and criteria related to the STS Path Trace diagnostics criteria, to be consistent with the new TIM-P text and criteria in Section 6.2.1.1.9 (see Section 6.2.3.2.3.A).
- Section 8, *SONET Operations Communications*, changes included:
 - The replacement of "X.25 DCN" terminology with "WAN" terminology throughout Section 8, and the corresponding removal of out-dated information (e.g., the former Section 8.1.6)
 - The revision of the requirement and text on the support of ES-IS at an OS/NE X.25 interface (ILR Issue ID 253-95, see Section 8.3.3.1)
 - The revision of several criteria related to the craftsperson/workstation interface so that they better reflect currently available technology (ILR Issue ID 253-132, see Sections 8.6.1 and 8.6.2)
 - The removal of the TARP Sequence Number (tar-seq) as a provisionable TARP PDU Field (ILR Issue ID 253-133, see Section 8.7.6)
- Appendix A, *Requirement-Object List*, was changed to reflect the new and revised criteria
- Appendix C, *SONET Operations Communications Lower Layers Protocol Profile* changes included:
 - The addition of the ENE, INE and GNE symbols to the SONET lower layer profile symbols list in Section C.1.2.2
 - The addition of the ENE, INE and GNE symbols to the appropriate tables in Sections C.5.2 and C.5.3
 - The addition of information regarding the status and profile columns of the tables in Section C.8
 - The removal of the TARP Sequence Number (tar-seq) as a provisionable TARP PDU Field (ILR Issue ID 253-133, see Section C.8.5.3).
- The *References* section was updated
- The *Requirement-Objects Index* section was updated.

2. Network Compatibility

This section presents criteria that are intended to help ensure SONET equipment compatibility with the existing network. These criteria deal with interfaces to the existing network and system performance, and are, for the most part, specific to SONET. Generic criteria applicable to SONET and other transmission systems (e.g., asynchronous fiber optic systems or digital radio systems) are found in GR-499-CORE, *Transport Systems Generic Requirements (TSGR): Common Requirements*.

This GR uses the terms Line, Section, and Path to delineate various paths of the transmission network that interconnect the SONET NEs. Figures 2-1 and 2-2 illustrate these transmission segments between various SONET NEs defined in the Glossary. Specific requirements for SONET NEs such as Add/Drop Multiplex (ADM), Terminal Multiplex (TM), Digital Cross-connect (DCS), and Regenerator (RGTR), are contained in referenced material.

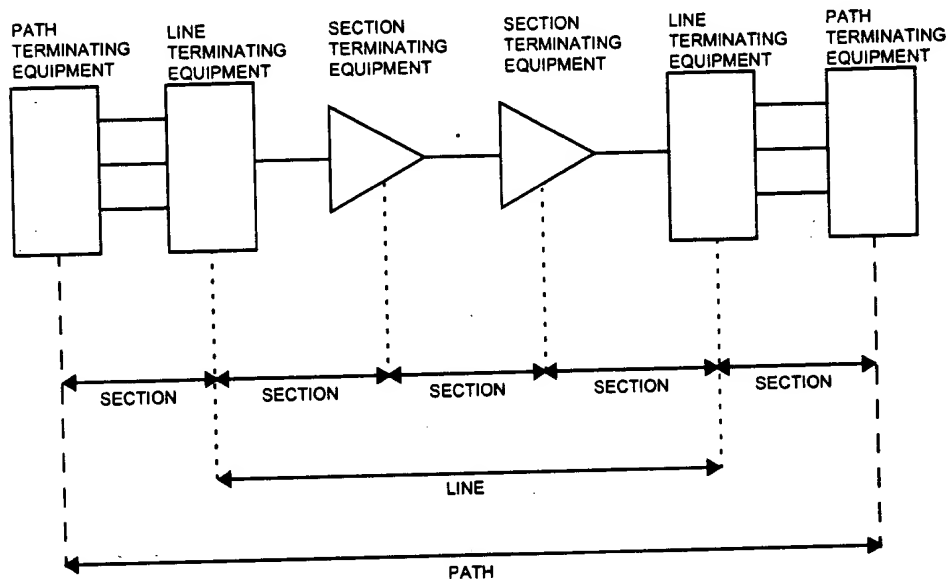


Figure 2-1. Simplified Diagram Depicting SONET Section, Line, and Path Definitions

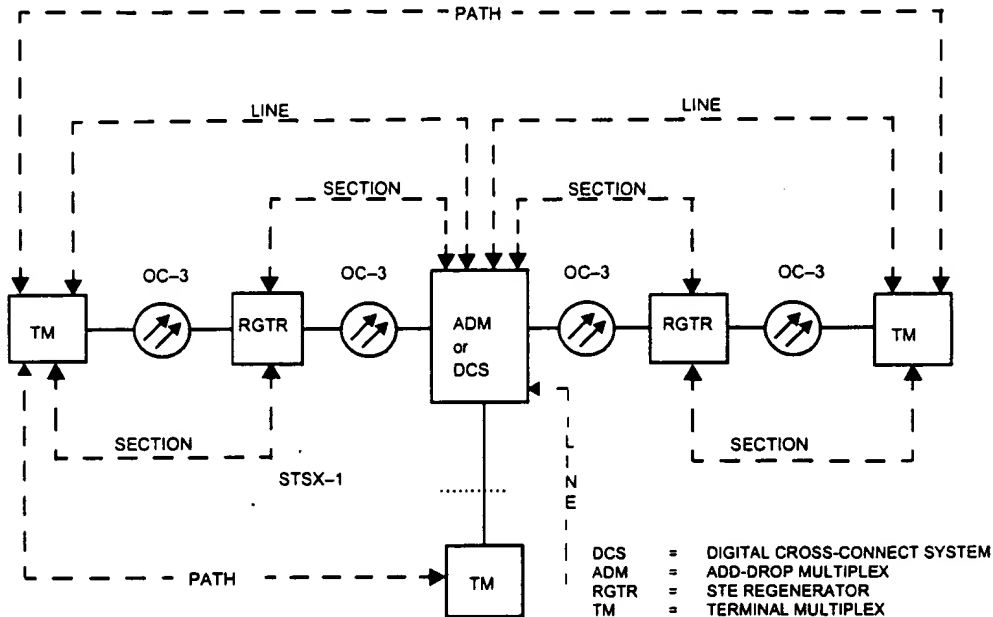


Figure 2-2. Diagram Illustrating SONET Section, Line, and Path Definitions

2.1 Network Element (NE) Interfaces

SONET NEs interface with existing telecommunications equipment such as other transmission systems, outside-plant fiber, operations systems, and power systems. This section discusses these interfaces.

2.1.1 Digital Signal Cross-Connect (DSX) Interface

This GR considers SONET fiber optic equipment that interfaces with other transmission equipment at electrical cross-connects. Interface requirements for interconnecting at the standard hierarchical cross-connects (DSX-1, DSX-1C, DSX-2, DSX-3, and DSX-4NA) are found in GR-499-CORE, and are based on ANSI T1.102, *Digital Hierarchy-Electrical Interfaces*, and ANSI T1.107, *Digital Hierarchy-Formats Specifications*. SONET NEs may also provide SONET electrical interfaces, and interconnect with other SONET NEs at SONET cross-connects (i.e., STSX-1s and STSX-3s) as described in Section 4.4.

- CR2-1** [1] SONET NEs may be required to provide electrical cross-connect facilities with patching and monitoring jacks for restoration and rearrangements.
- CR2-2** [2] A bridging repeater or equivalent may be required to allow rearrangement from the monitor jack.

2.1.1.1 Electrical Cable Distance

The maximum cable distance between the terminal and the cross-connect frame is a function of the interface signal hierarchical level. For any DS1, DS1C, DS2, and DS3 interfaces provided, the cross-connect interface requirements of GR-499-CORE apply under the reference cable and cable length combinations listed in Table 2-1.

Table 2-1. Reference Cable Lengths for Testing DS_n Interfaces

Rate	Cable Length on Each Side of DSX	Reference Cable Assumed ^a
DS1	655 ft (199.6 m)	AT&T Technologies, Inc. 22ga. ABAM (or equivalent)
DS1C	655 ft (199.6 m)	AT&T Technologies, Inc. 22ga. ABAM (or equivalent)
DS2	1000 ft (304.8 m)	AT&T Technologies, Inc. 22ga. ABAM (or equivalent)
DS3	450 ft (137.2 m)	AT&T Technologies, Inc. 728A Coaxial (or equivalent)

- a. Each length and cable-type pair listed is for testing purposes. A user may specify an equivalent length and cable-type combination for a particular application.

2.1.1.2 Maintenance Signal Compatibility

To prevent unwanted propagation of alarms and to help sectionalize failures in a manner consistent with the existing networks, SONET NEs provide and interact with maintenance signals at the standard hierarchical rates. See Section 6.2.1.

2.1.2 Interface to Fiber Distributing Frame and Optical DSXs

Fiber distributing frames and optical DSXs allow for connecting fiber cable without disassembling splices. Information on fiber distributing frames and optical DSXs is in GR-449-CORE, *Generic Requirements and Design Considerations for Fiber Distributing Frames*, and TA-NPL-000464, *Generic Requirements and Design Considerations for Optical Digital Signal Cross-Connect Systems*, respectively.

- R2-3** [3] Suppliers shall provide a description of their optical fiber distributing frame or optical DSXs. This description shall include the number of terminations, storage provisions for excess fiber, and the type of fiber cable connectors provided (see Section 4).

2.1.3 Operations Systems Interface

The protocols, languages, and content of the generic networking requirements for interfaces with BCC operations systems are described in Section 8. The procuring BCC determines the interface arrangement.

2.1.4 Synchronization Interface

SONET uses the existing synchronization network as described in GR-436-CORE, *Digital Network Synchronization Plan* and ANSI T1.101, *Synchronization Interface Standards for Digital Networks*. The goal is to create a fully synchronous optical hierarchy by ensuring that all SONET NEs derive timing traceable to a primary reference source.

SR-NWT-002224, *SONET Synchronization Planning Guidelines*, provides many details about the integration of SONET NEs into the BCCs' synchronization networks. A SONET network may use more than one primary reference source. A primary reference source is equipment that provides a timing signal whose long-term accuracy is maintained at 10^{-11} or better with verification to Universal Time Coordinated (UTC), and whose timing signal is used as the basis of reference for the control of other clocks within a network.

The clocks used to synchronize SONET NEs are stratum 3 (or better quality) clocks as described in ANSI T1.101. Thus, the timing for SONET signals is normally traceable to a primary reference source as GR-436-CORE and ANSI T1.101 describe.

Requirements for the short-term stability of the SONET signals are described in Section 5.4, which provides detailed criteria on the SONET NE synchronization interface, timing, and clocks.

2.1.5 Power

The BCC provides dc power to the supplier's equipment with a nominal voltage of -48 V. Requirements on SONET equipment pertaining to voltage limits, electrical noise, and current drain are in GR-499-CORE.

2.2 End-to-End Performance Criteria

This section addresses issues related to SONET system performance from the standpoint of compatibility with the existing network. These criteria ensure that a SONET transport system performs as well as other systems that transport the standard asynchronous hierarchical signals (DS1, DS1C, DS2, and DS3).

The subjects covered in the following subsections include availability (i.e., maximum downtime), error performance, protection switching performance, jitter, and transmission delay.

2.2.1 Availability and Reliability

Service availability and system reliability are critical issues for the BCCs as well as their customers. Equipment suppliers should consider reliability issues during all phases of system design and manufacture. In terms of design goals, the equipment supplier needs to understand the impact of individual equipment reliability on end-to-end service availability.

The underlying foundation for many Bellcore system reliability criteria is an end-to-end two-way availability objective of 99.98% for interoffice applications (0.02% unavailability or 105 min/yr maximum downtime). The corresponding objective for loop transport between the central office and the customer's premises is 99.99%. In interoffice transport, the objective refers to a two-way broadband channel (e.g., DS3, STS-N, or OC-N) over a 250-mile path; in loop applications, the objective applies to a two-way narrowband channel (e.g., DS0 or equivalent). In either case, the objective is meant as a long-term average over a large area for regular (non-special) service offerings, such as "plain old telephone service" for loop customers.¹

Starting with these end-to-end objectives, a top-down approach allocates maximum downtime to various parts of a generic model network, described in terms of a "hypothetical reference circuit". SONET systems present a challenging problem because a BCC can eventually "mix and match" equipment from different suppliers. To ensure that the total downtime does not exceed the end-to-end objective, this situation for SONET forces allocations down to the level of individual NEs. Special consideration is needed for rings and for metropolitan applications.

A comprehensive discussion of SONET availability, including specific criteria and associated rationale, is presented in TR-NWT-000418, *Generic Reliability Assurance Requirements for Fiber Optic Transport Systems*. Specific availability criteria for SONET NEs are provided in Section 2 of that TR.

1. Different levels of performance could be requested by a customer or assured by a BCC, based on matters outside the general discussion here.

2.2.2 Protection Switching Performance

APS systems increase system availability by automatically substituting a protection line (including optical transmitters and receivers) for a failed line. Digital channel banks and other digital NEs produce false signaling states and eventually a Carrier Group Alarm (CGA) after receiving a signal with sufficiently high Bit Error Ratio (BER) or incoming signal failure for a specified time. Thus, a protection switching system must recognize both situations and respond in a sufficiently short time. Section 5 of GR-499-CORE contains generic criteria on protection switching for fiber optic systems.

The physical performance of line protection switching is characterized by the time to detect certain switching thresholds (based on BER) and the time to physically complete the switch. Section 5.3 of this document presents criteria for SONET protection switching.

2.2.3 Error Performance

Error performance criteria are concerned with such parameters as BER and Errored Seconds observed at standard hierarchical interface rates (DS1, DS1C, DS2, and DS3). Section 4 of GR-499-CORE gives definitions and requirements for these performance parameters.

2.2.4 Jitter

Timing jitter may arise from a number of sources within the digital network. Timing jitter is defined as the short-term variations of the significant instants of a digital signal from their ideal positions in time, where short-term implies phase oscillations of frequencies greater than some demarcation frequency. Section 5.6 contains the jitter criteria applicable to SONET NEs and systems.

2.2.5 Transmission Delay

Guidelines for transmission delay in an exchange access network are given in ANSI T1.506A-1992, *Telecommunications - Network Performance - Specifications for Switched Exchange Network (Absolute Round-Trip Delay)*, ANSI T1.508-1992, *Telecommunications - Network Performance - Loss Plan for Evolving Digital Networks*, and ANSI T1.508A-1993, *Telecommunications - Network Performance - Loss Plan for Evolving Digital Networks*. Specific SONET NEs have specific transmission delay criteria, as described in the appropriate NE-specific GRs, TRs and TAs. When incorporating SONET NEs into a transmission path, the processing and propagation delays contributed by the NEs and the propagation delay contributed by the interconnecting media must be taken into consideration in the context of the guidelines in the above ANSI standards.

3. Rates and Formats

This section defines the rates and formats for SONET signals. A primary goal in defining these signals is to articulate a synchronous hierarchy that has sufficient flexibility to carry many different capacity signals. This is realized by defining a basic module with a bit rate of 51.840 Mb/s, and a byte-interleaved multiplex scheme that results in a family of signals with rates of N times 51.840 Mb/s, where N is an integer (Section 3.1).

The basic module can be divided into a portion assigned to overhead and a portion that carries the payload (Section 3.2). This payload portion can be used to transport DS3 signals or a variety of sub-DS3 signals. Because some signals requiring transport have rates greater than the basic rate [e.g., some Broadband Integrated Services Digital Network (B-ISDN) applications], a technique of linking several basic modules together to build a transport signal of increased capacity is described (Section 3.2.3). To maintain a consistent payload structure while providing for the transport of a variety of lower rate payloads (e.g., DS1, DS1C, and DS2 signals), a structure called the Virtual Tributary (VT) is defined (Section 3.2.4). Payloads below the DS3 rate are transported within a VT structure.

Different types of overhead are defined for functions including maintenance, protection switching, frequency justification, orderwire, identification, and user channels (Section 3.3). Also, growth channels are identified to allow for future uses not defined or conceived of at this time. A layered approach to overhead is established, whereby overhead bandwidth is allocated to a layer based on the function addressed by that particular layer. This layered approach allows the creation of equipment that is not required to access all layers of overhead, thereby allowing equipment implementations to meet different needs.

Section 3.4 considers the mapping of various payloads into payload envelopes. Section 3.5 describes payload pointers, which are mechanisms that allow the payload envelopes to slide relative to the overhead, thus permitting accommodation of different signal phases and frame rates in multiplexing.

3.1 Synchronous Hierarchical Rates

The Synchronous Transport Signal-level 1 (STS-1) is the basic module in SONET. It has a bit rate of 51.840 Mb/s. The optical counterpart of the STS-1 is the Optical Carrier - level 1 (OC-1) signal, and the electrical counterpart of the STS-1 is the STS-1 electrical [or Electrical Carrier - level 1 (EC-1)] signal defined in Section 4.4.

The definition of the first level also defines the entire hierarchy of SONET signals because higher-level signals are obtained by synchronously multiplexing lower-level modules. When lower-level modules are multiplexed together, the result is denoted as an STS- N (where N is an integer), which can then be converted to an OC- N or STS- N electrical signal. There is an integer multiple relationship between the rates of the basic STS-1 module and the OC- N or STS- N electrical signals (i.e., the rate of an OC- N is equal to N times the rate of an STS-1).

SONET systems support only certain values of N. Table 3-1 lists these values for the standard STS-N electrical and OC-N interface signals up through N = 48, along with the corresponding line rates. Values of N greater than 48 are addressed in other documents (e.g., GR-1377-CORE, *SONET OC-192 Transport System Generic Criteria*).

Table 3-1. Line Rates for Standard SONET Interface Signals (N ≤ 48)

OC-N Level	STS-N Electrical Level	Line Rate (Mb/s)
OC-1	STS-1 electrical	51.84
OC-3	STS-3 electrical	155.52
OC-12	—	622.08
OC-24	—	1244.16
OC-48	—	2488.32

3.2 Transport Format

The SONET transport format presented here is based on ANSI T1.105, *Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates and Formats*. The format definition in the following sections designates some bits and bytes as undefined. Suppliers are likely to introduce enhanced features by using these bits and bytes in a nonstandard manner. Network providers who wish to deploy these nonstandard features should study the network implications jointly with the supplier, and recognize the potential equipment incompatibilities.

- R3-1** [4] A SONET NE shall have the capability to ignore the values contained in all undefined and unused bits and bytes [except for Bit Interleaved Parity (BIP)-8 calculations] to prevent misinterpretation of the received patterns.

Undefined bits and bytes are those for which no standard use has been defined for the transmitting NE. Unused bit and bytes are those for which no standard use has been defined for the NE when it receives the SONET signal. (The criteria on the use of the currently defined overhead bits and bytes are summarized in Table 5-2.)

- O3-2** [5] A SONET NE should send all-zeros patterns (before scrambling) in undefined bits and bytes. All-zeros patterns should also be sent in defined bits and bytes if the NE does not support the defined function or if the function has been disabled by the user.

Note that for many bits and bytes, O3-2 [5] is overridden by the AIS generation requirements in Section 6.2.1.2 when the NE is transmitting AIS. Also, as discussed in various sections of this document, an NE that transmits particular non-zero values in certain undefined bits and bytes can still meet the objective. For example, although the B1 byte position is not defined for a drop-side STS-1 electrical signal, an NE may transmit either all-zeros or the Section BIP-8 code in that byte (see Section 3.3.2.1).

- R3-3** [6] If a supplier introduces a nonstandard feature employing SONET overhead, the supplier shall disclose such use of overhead, and furnish the network provider with an equipment option to disable the feature (including the transmission of the nonstandard messages).

3.2.1 Frame Structure of the STS-1

An STS-1 is a specific sequence of 810 bytes (6480 bits), which includes various overhead bytes and an envelope capacity for transporting payloads. It can be depicted as a 90 column by 9 row structure, as shown in Figure 3-1. With a frame length of 125 μ s (i.e., 8000 frames per second), the STS-1 has a bit rate of 51.840 Mb/s. Using the structure in Figure 3-1, the order of transmission of bytes is row-by-row, from left to right.

- R3-4** [7] The structure of an STS-1 shall be as shown in Figure 3-1.¹

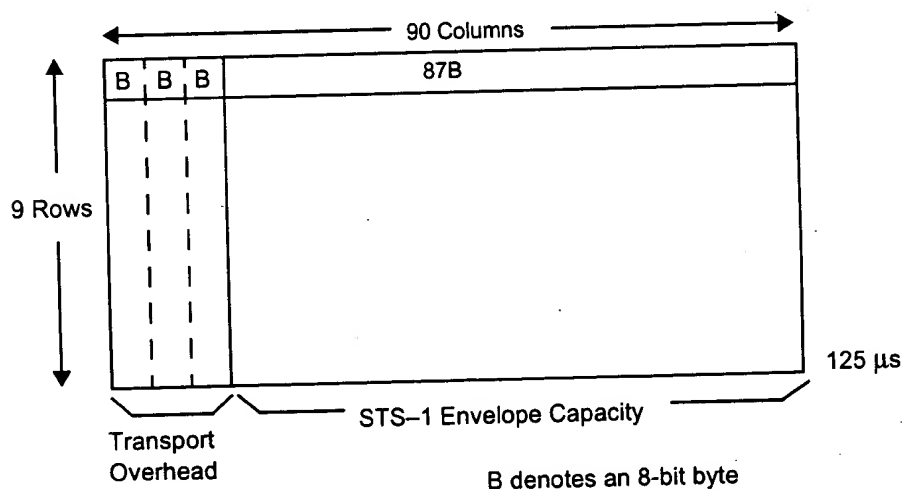


Figure 3-1. STS-1 Frame

1. What is meant by "the structure of an STS-1 shall be as shown ..." is that the byte sequence shall be such that it can be mapped into the frame structure shown.

- R3-5** [8] In each byte of the STS-1, the most-significant bit shall be transmitted first, as shown in Figure 3-2.

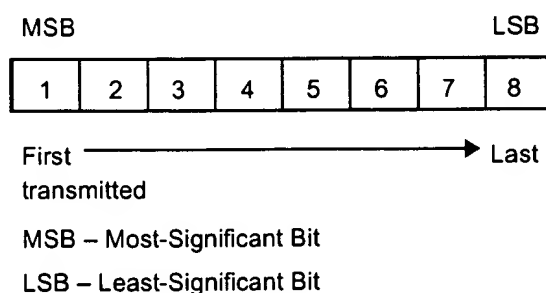


Figure 3-2. Bit Position Numbering

3.2.1.1 Transport Overhead

As Figure 3-1 shows, the first three columns of the STS-1 frame are the Transport Overhead. These three columns contain 27 bytes, of which nine bytes are overhead for the Section layer (i.e., Section Overhead), and 18 bytes are overhead for the Line layer (i.e., Line Overhead). Section 3.3 contains the details of these overhead allocations. The remaining 87 columns constitute the STS-1 Envelope Capacity.

3.2.1.2 STS-1 Envelope Capacity and Synchronous Payload Envelope (SPE)

Figures 3-3, 3-4, and 3-5 depict the STS-1 SPE, which occupies the STS-1 Envelope Capacity. The STS-1 SPE consists of 783 bytes, and can be depicted as an 87 column by 9 row structure. Column 1 contains nine bytes, designated as the STS Path Overhead (POH). Two columns (columns 30 and 59) are not used for payload, but are designated as the “fixed stuff” columns. The 756 bytes in the remaining 84 columns are designated as the STS-1 Payload Capacity.

- R3-6** [9] The structure of an STS-1 SPE shall be as shown in Figure 3-4.

The bytes in the fixed stuff columns are undefined, so the objective in Section 3.2 (to set them to all-zeros) is applicable. However, several possible uses for those bytes have been discussed in the standards bodies, and some suppliers may choose to use them for proprietary purposes. Therefore, for compatibility between the STS-1 Path BIP-8 calculation in SONET (which covers all 87 columns of the STS-1 SPE) and the VC-3 Path BIP-8 calculation in the international Synchronous Digital Hierarchy (SDH) in ITU-T

Recommendation G.709, *Synchronous multiplexing structure* (which covers all 85 columns of the VC-3), the following requirement has been added:

- R3-7** [10] The values used to stuff columns 30 and 59 of each STS-1 SPE shall produce even parity in the calculation of the STS-1 Path BIP-8 (see Section 3.3.2.3).

The STS-1 SPE may begin anywhere in the STS-1 Envelope Capacity. Typically, it begins in one STS-1 frame and ends in the next (although it may be wholly contained in one frame). The STS Payload Pointer contained in the Transport Overhead designates the location of the byte where the STS-1 SPE begins. Section 3.5.1 describes STS Payload Pointers.

STS POH is associated with each payload and is used to communicate various information from the point where a payload is mapped into the STS-1 SPE to where it is delivered. Section 3.3.2.3 contains details on the STS POH.

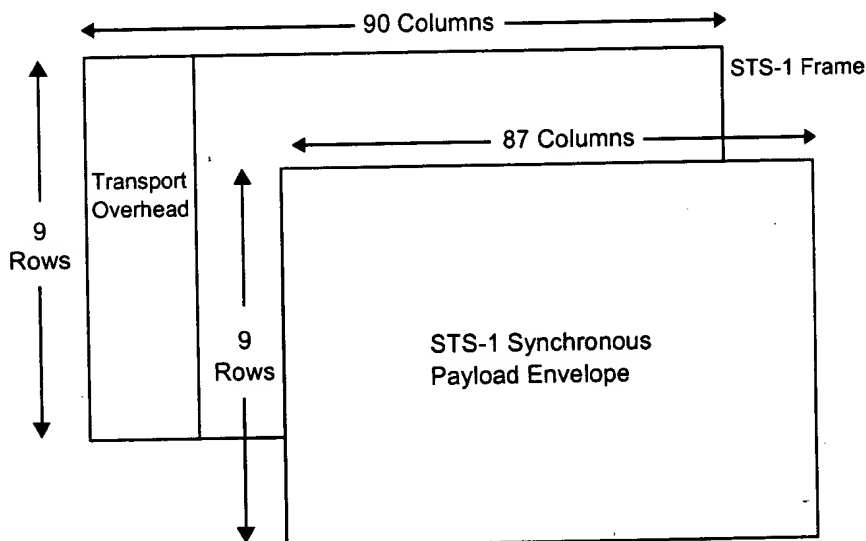


Figure 3-3. STS-1 Synchronous Payload Envelope

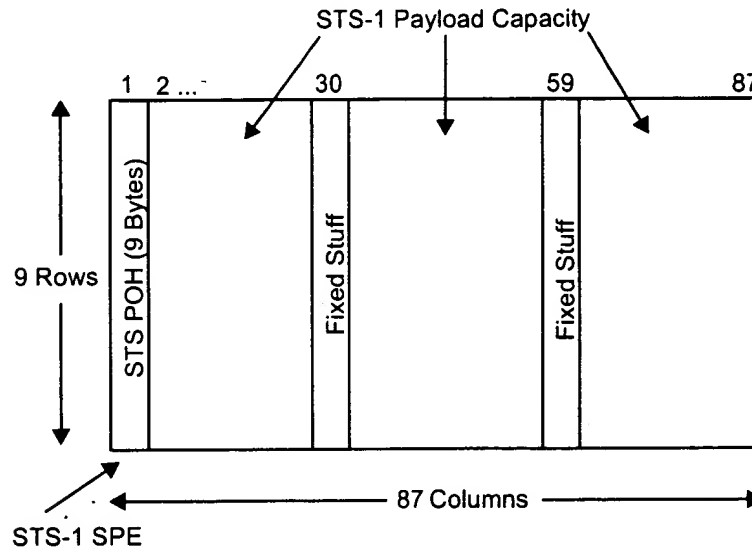


Figure 3-4. STS-1 SPE with STS-1 POH and STS-1 Payload Capacity Illustrated

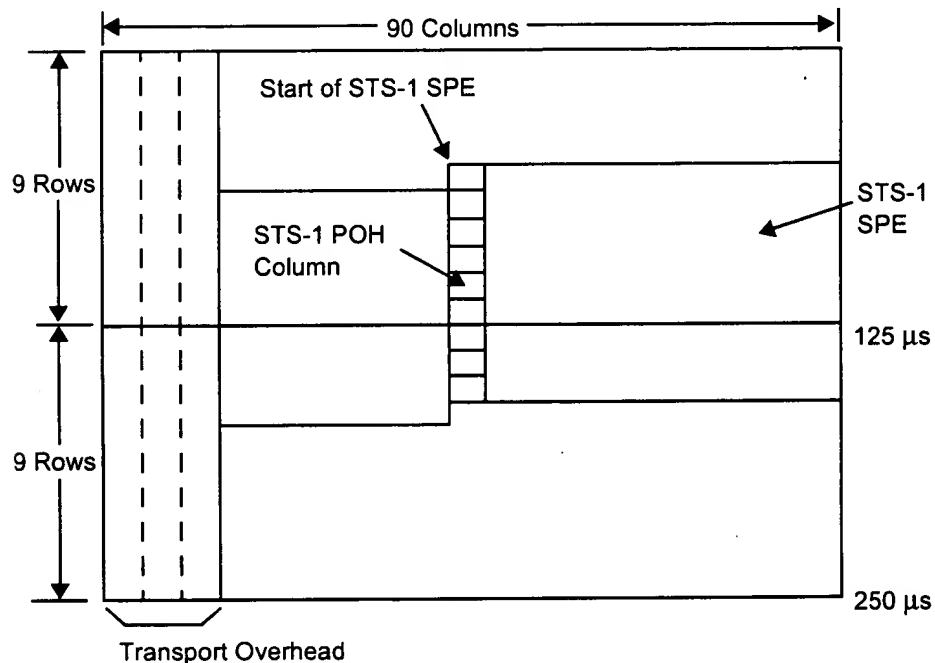


Figure 3-5. STS-1 SPE in Interior of STS-1 Frames

3.2.2 Frame Structure of the STS-N

An STS-N is a specific sequence of $N \times 810$ bytes that can be depicted as the structure shown in Figure 3-6. The STS-N is formed by byte-interleaving STS-1 and STS-M ($3 \leq M < N$) modules. The Transport Overhead of the individual STS-1 and STS-M modules are frame aligned before interleaving, but the associated STS SPEs are not required to be aligned because each STS-1 has a Payload Pointer to indicate the location of the SPE (or to indicate concatenation). Section 5.1.1 contains the interleaving requirements.

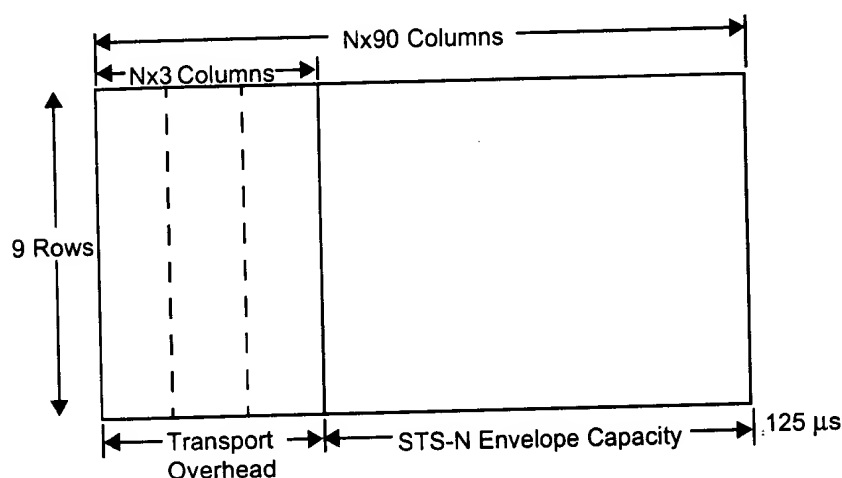


Figure 3-6. STS-N Frame

3.2.3 STS Concatenation

Multiple STS-1 SPEs are needed to transport Super Rate payloads, such as some B-ISDN ATM payloads. To accommodate such a payload, an STS-Nc module is formed by linking N constituent STS-1s together in fixed phase alignment. The Super Rate payload is then mapped into the resulting STS-Nc SPE for transport. The STS-Nc SPE can be carried by an OC-N, STS-N electrical, or higher rate signal.

The need for a SONET NE to be able to generate, multiplex, switch, transport, or terminate STS-Nc SPEs depends on the functionality of that NE. Criteria for specific types of NEs are contained in the individual SONET NE criteria documents (e.g., GR-496-CORE, *SONET Add-Drop Multiplexer (SONET ADM) Generic Criteria*).

- R3-8** [11] If an NE supports the multiplexing, switching, or transport of STS-Nc SPEs, then it shall treat each STS-Nc SPE as a single entity.

Concatenation indicators contained in the second through Nth STS Payload Pointers are used to show that the STS-1s of an STS-Nc are linked together.

An STS-Nc SPE consists of $N \times 783$ bytes, and can be depicted as an $N \times 87$ column by 9 row structure, as shown in Figure 3-7 (which also shows the STS-Nc Payload Capacity). Only one set of STS POH is required in the STS-Nc SPE. The STS-Nc SPE is carried within the STS-Nc so that the STS POH always appears in the first of the N STS-1s that make up the STS-Nc.

In all of the Super Rate payload mappings contained in this document, the first $(N/3)-1$ columns of the STS-Nc SPE following the STS POH are not used for payload, but are designated as fixed stuff columns, (i.e., columns of undefined bytes, see Section 3.2).² Since the STS-Nc SPE is treated as a single entity, the presence or absence of fixed stuff columns only affects the equipment that generates and terminates the SPE. Therefore, future payload mappings could be defined where these columns are used in the payload mapping.

To date, only mappings into STS-3c and STS-12c SPEs have been defined in this document (see Section 3.4.3). Other mappings requiring different values of N could possibly be defined in the future; however, many SONET products may support the transport of only certain size STS-Nc SPEs (i.e., STS-3c, STS-12c and STS-48c SPEs). Therefore, an STS-Nc path terminating product's use of other values of N may restrict its applicability in multi-product configurations.

R3-9 [12] The structure of an STS-Nc SPE shall be as shown in Figure 3-7.

Sections 3.5.1.4 and 5.1.2 provide additional details on STS concatenation: Section 3.3.2 describes the overhead assignments, and Figure 3-8 illustrates the transport overhead assignments for an STS-3 electrical or OC-3 signal carrying an STS-3c SPE. Section 3.4.3 describes mappings of Super Rate payloads into STS-Nc SPEs.

2. Note that this implies N must be divisible by 3. It also means that an STS-3c SPE has no fixed stuff columns (although the payload mapping contained in the SPE may contain one or more columns of fixed stuff bytes).

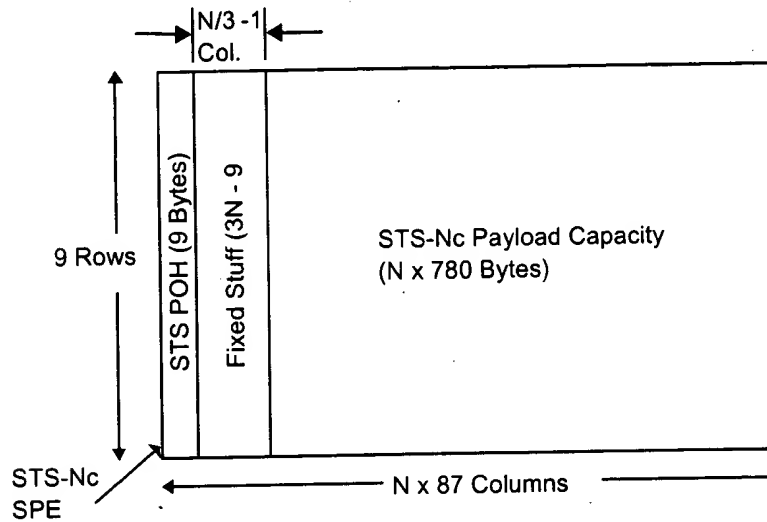


Figure 3-7. STS-Nc SPE

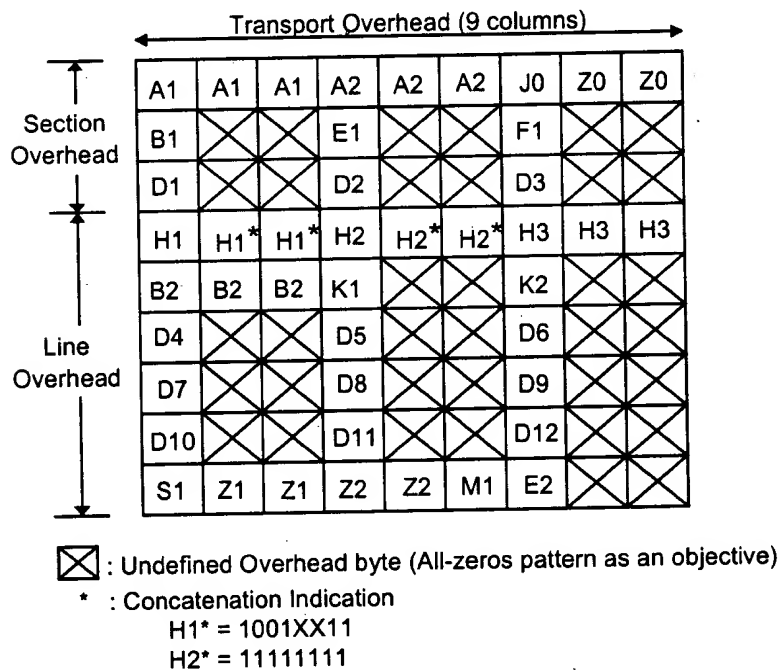


Figure 3-8. Transport Overhead Assignment, OC-3 Carrying an STS-3c SPE

3.2.4 Virtual Tributary (VT) Structure

The VT structure is designed for transport and switching of sub-ST_S-1 rate payloads. There are four sizes of VTs: VT1.5 (1.728 Mb/s), VT2 (2.304 Mb/s), VT3 (3.456 Mb/s), and VT6 (6.912 Mb/s). These are illustrated in Figure 3-9. In the 87-column by 9-row structure of the ST_S-1 SPE, these VTs occupy 3, 4, 6, and 12 columns, respectively.

To accommodate a mix of VT sizes efficiently, the VT-structured ST_S-1 SPE is divided into seven VT groups. Each VT group occupies 12 columns of the 87-column ST_S-1 SPE, and may contain 4 VT1.5s, 3 VT2s, 2 VT3s, or 1 VT6. A VT group can contain only one size of VTs; however, a different VT size is allowed for each VT group in an ST_S-1 SPE.

Figures 3-10, 3-12, 3-14, and 3-16 each show all seven VT groups in an ST_S-1 SPE containing one of the four VT sizes. The tables in Figures 3-11, 3-13, 3-15, and 3-17 define the relationship between the VT group number and VT number, and the columns in the ST_S-1 SPE from Figures 3-10, 3-12, 3-14, and 3-16, respectively. These tables are applicable in all cases, including VT groups in an ST_S-1 SPE with different VT sizes. Figures 3-18 and 3-19 illustrate an example where the first four VT groups contain VT1.5s, VT2s, VT3s, and VT6s.

- R3-10** [13] The structure of a VT-structured ST_S-1 SPE shall be consistent with the structures shown in Figures 3-9 through 3-19.

In addition to the division of VTs into VT groups, a 500-μs structure called a VT Superframe is defined for each VT. The VT Superframe contains the V1 and V2 bytes (the VT Payload Pointer), the V3 byte (the VT pointer action byte), the V4 byte (an undefined byte), and the VT Envelope Capacity, which in turn contains the VT SPE. The VT Envelope Capacity, and therefore the size of the VT SPE, is different for each VT size, as shown in Figure 3-20. V1 is the first byte in the VT Superframe, while V2 through V4 appear as the first bytes in the following frames of the VT Superframe, regardless of the VT size.

- R3-11** [14] Four consecutive 125-μs frames of the VT-structured ST_S-1 SPE shall be organized into a 500-μs superframe, the phase of which is indicated by the H4 (Indicator) byte in the STS POH (see Section 3.4.1).

The VT Payload Pointer provides for flexible and dynamic alignment of the VT SPE within the VT Envelope Capacity, independent of other VT SPEs. Section 3.5.2 further describes the VT Payload Pointers. Figure 3-21 illustrates the VT SPEs corresponding to the four VT sizes. Each VT SPE contains four bytes of VT POH (V5, J2, Z6, and Z7), and the remaining bytes constitute the VT Payload Capacity, which is different for each VT size. Section 3.4.1 describes the mappings for various payloads (e.g., DS1, DS1C, DS2) into VT SPEs.

- R3-12** [15] The structure of a VT SPE shall be as shown in Figure 3-21.

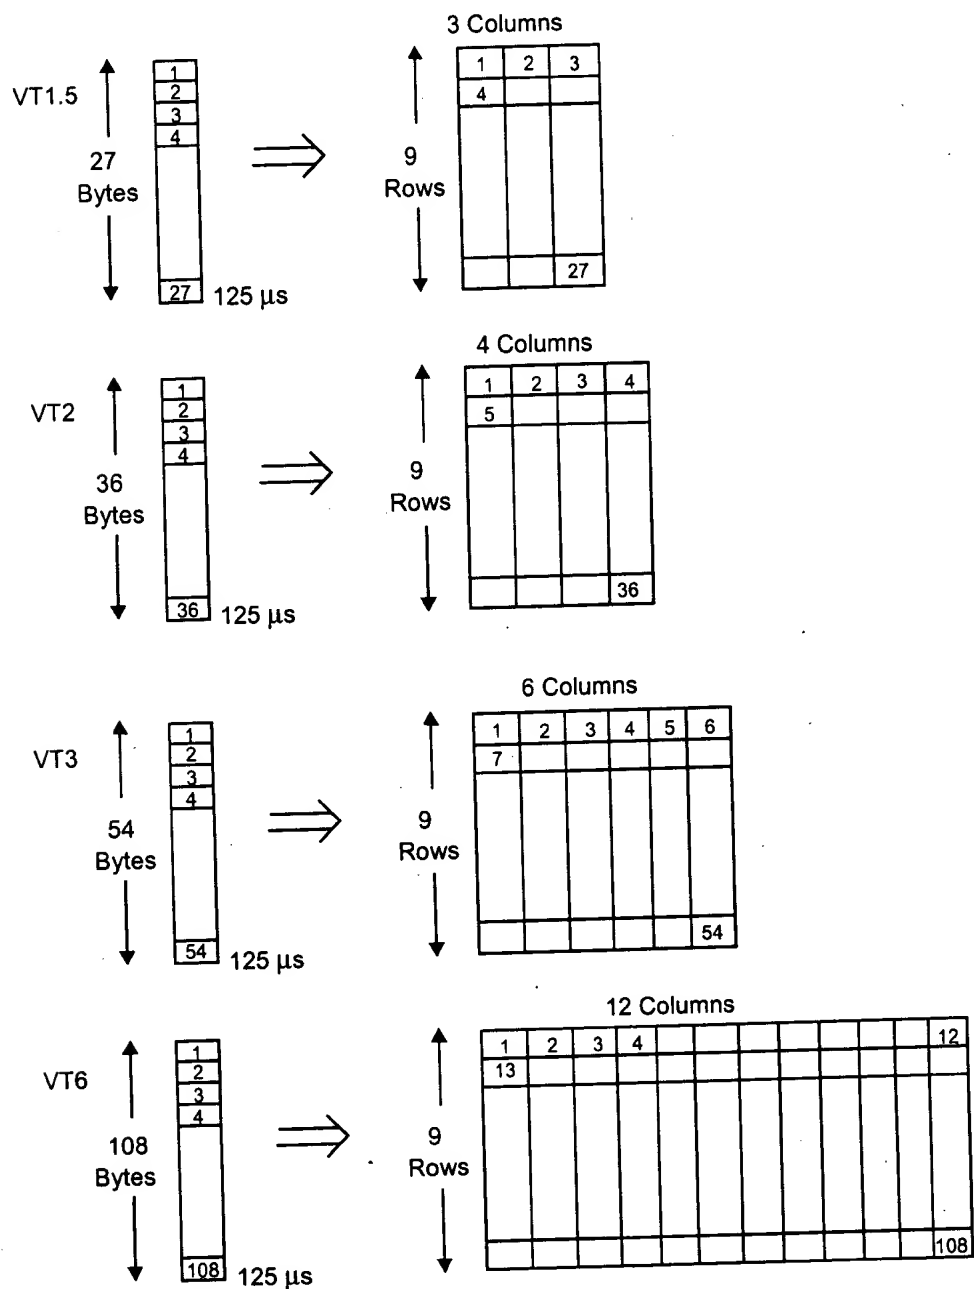


Figure 3-9. VT Sizes

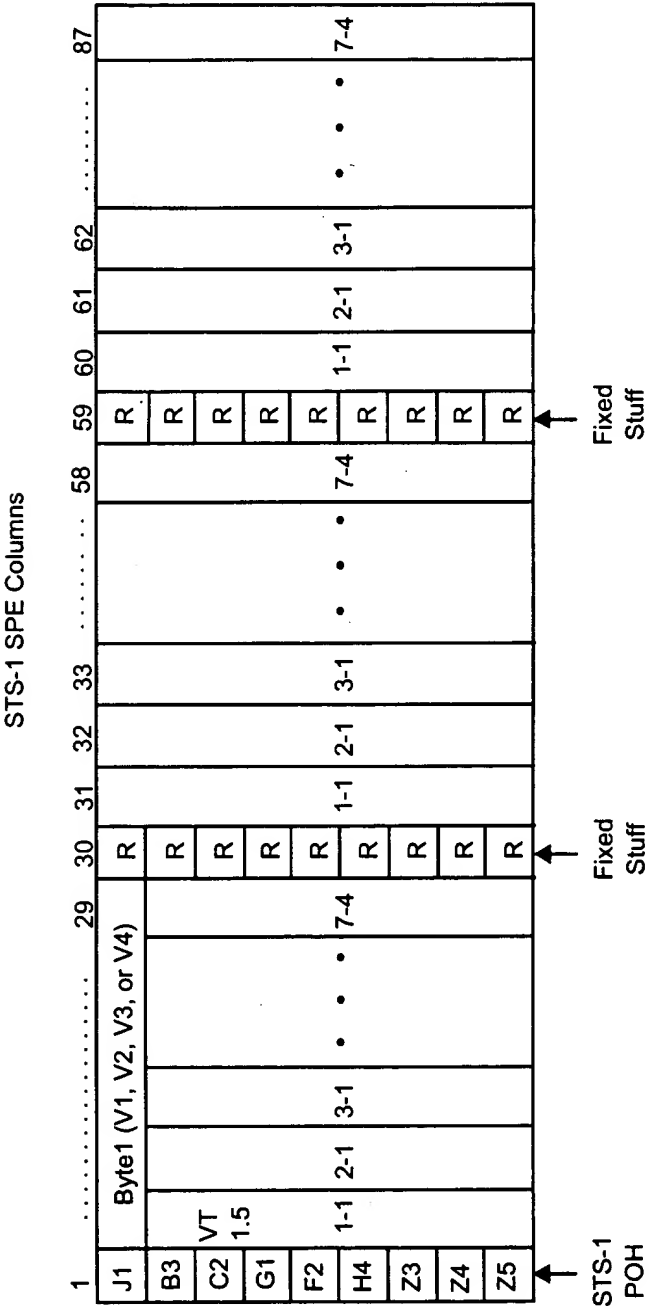


Figure 3-10. VT Structured STS-1 SPE: All VT1.5s

<u>VT Group #, VT #</u>	<u>Columns</u>	
1, 1	2, 31, 60	
2, 1	3, 32, 61	
3, 1	4, 33, 62	
4, 1	5, 34, 63	Column 1 = STS-1 POH
5, 1	6, 35, 64	30 = Fixed Stuff
6, 1	7, 36, 65	59 = Fixed Stuff
7, 1	8, 37, 66	
1, 2	9, 38, 67	
2, 2	10, 39, 68	
3, 2	11, 40, 69	
4, 2	12, 41, 70	
5, 2	13, 42, 71	
6, 2	14, 43, 72	
7, 2	15, 44, 73	
1, 3	16, 45, 74	
2, 3	17, 46, 75	
3, 3	18, 47, 76	
4, 3	19, 48, 77	
5, 3	20, 49, 78	
6, 3	21, 50, 79	
7, 3	22, 51, 80	
1, 4	23, 52, 81	
2, 4	24, 53, 82	
3, 4	25, 54, 83	
4, 4	26, 55, 84	
5, 4	27, 56, 85	
6, 4	28, 57, 86	
7, 4	29, 58, 87	

Figure 3-11. VT1.5 Locations

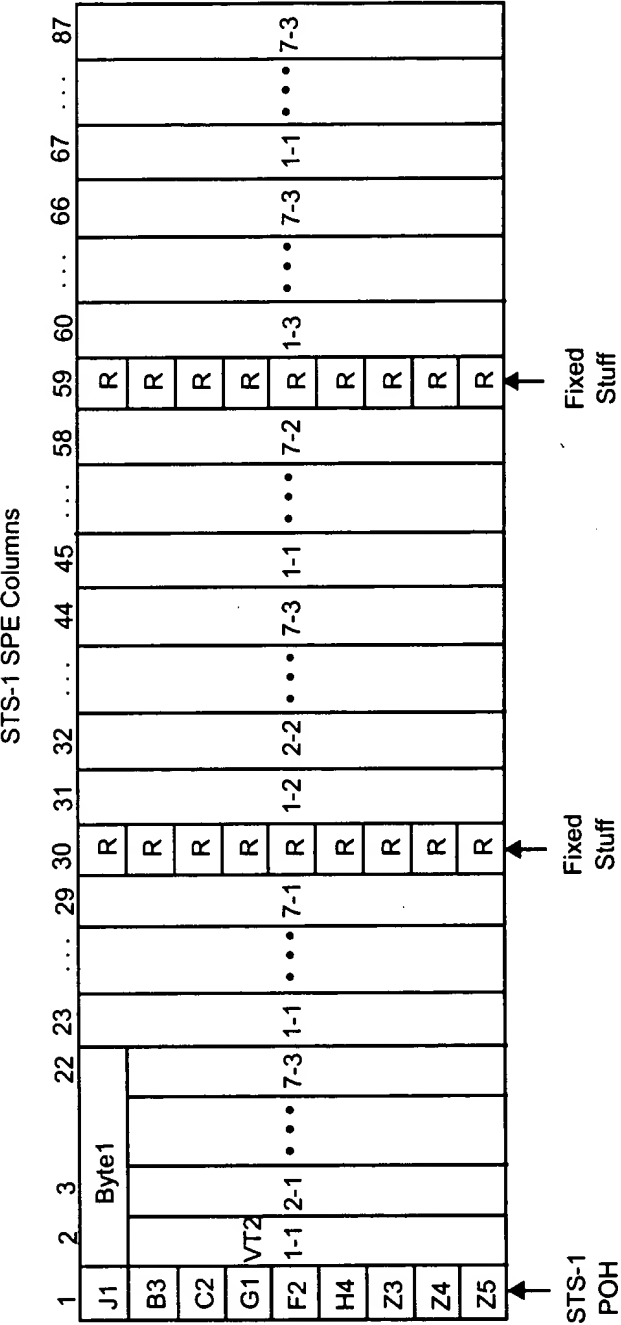


Figure 3-12. VT Structured STS-1 SPE: All VT2s

<u>VT Group #, VT #</u>	<u>Columns</u>	
1, 1	2, 23, 45, 67	
2, 1	3, 24, 46, 68	
3, 1	4, 25, 47, 69	
4, 1	5, 26, 48, 70	Column 1 = STS-1 POH
5, 1	6, 27, 49, 71	30 = Fixed Stuff
6, 1	7, 28, 50, 72	59 = Fixed Stuff
7, 1	8, 29, 51, 73	
1, 2	9, 31, 52, 74	
2, 2	10, 32, 53, 75	
3, 2	11, 33, 54, 76	
4, 2	12, 34, 55, 77	
5, 2	13, 35, 56, 78	
6, 2	14, 36, 57, 79	
7, 2	15, 37, 58, 80	
1, 3	16, 38, 60, 81	
2, 3	17, 39, 61, 82	
3, 3	18, 40, 62, 83	
4, 3	19, 41, 63, 84	
5, 3	20, 42, 64, 85	
6, 3	21, 43, 65, 86	
7, 3	22, 44, 66, 87	

Figure 3-13. VT2 Locations

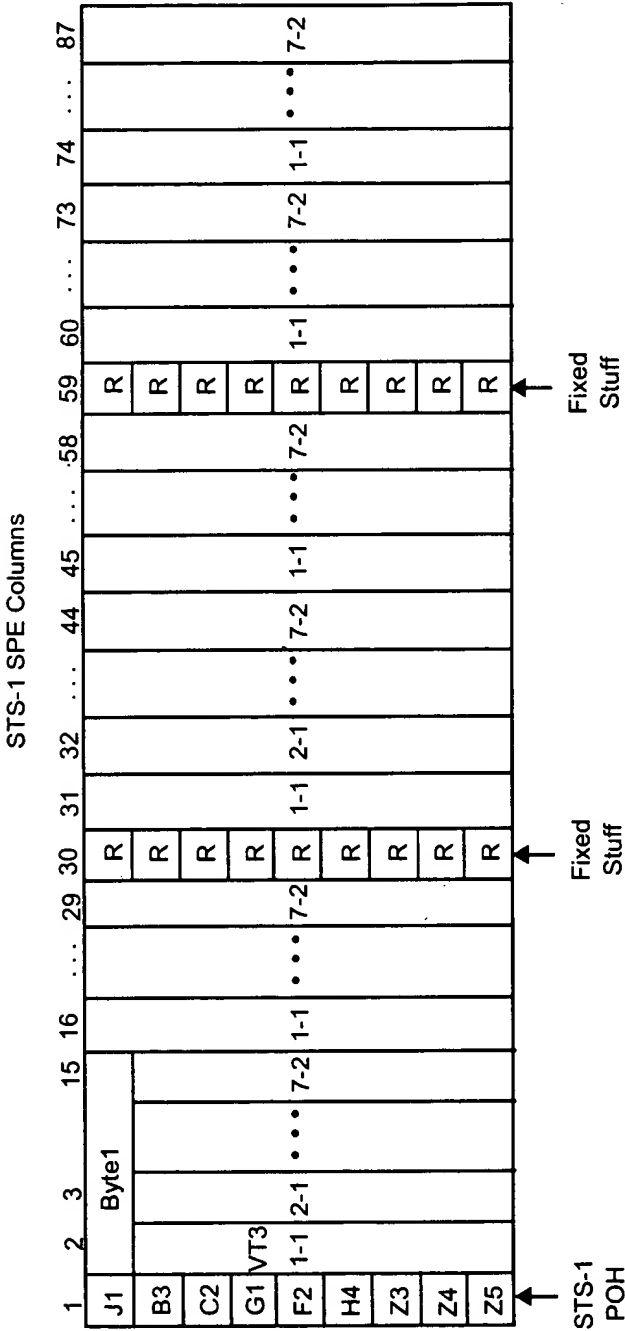


Figure 3-14. VT Structured STS-1 SPE: All VT3s

<u>VT Group #, VT #</u>	<u>Columns</u>	
1, 1	2, 16, 31, 45, 60, 74	
2, 1	3, 17, 32, 46, 61, 75	
3, 1	4, 18, 33, 47, 62, 76	
4, 1	5, 19, 34, 48, 63, 77	Column 1 = STS-1 POH
5, 1	6, 20, 35, 49, 64, 78	30 = Fixed Stuff
6, 1	7, 21, 36, 50, 65, 79	59 = Fixed Stuff
7, 1	8, 22, 37, 51, 66, 80	
1, 2	9, 23, 38, 52, 67, 81	
2, 2	10, 24, 39, 53, 68, 82	
3, 2	11, 25, 40, 54, 69, 83	
4, 2	12, 26, 41, 55, 70, 84	
5, 2	13, 27, 42, 56, 71, 85	
6, 2	14, 28, 43, 57, 72, 86	
7, 2	15, 29, 44, 58, 73, 87	

Figure 3-15. VT3 Locations

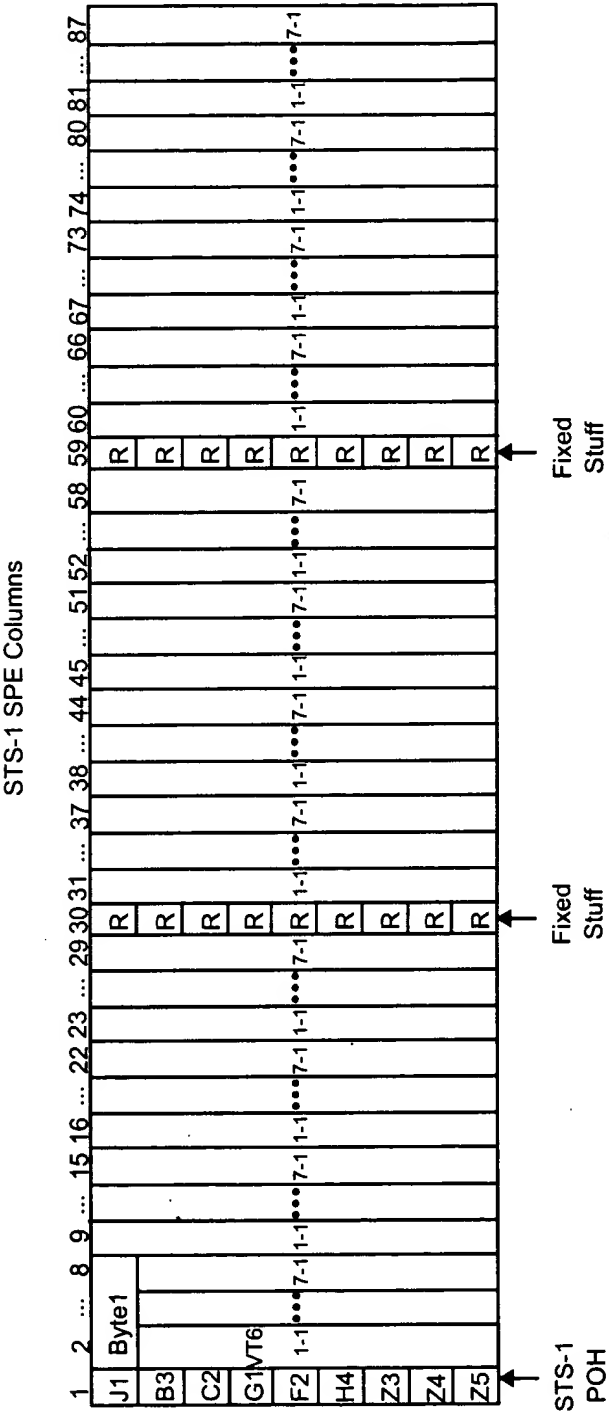


Figure 3-16. VT Structured STS-1 SPE: All VT6s

<u>VT Group #, VT #</u>	<u>Columns</u>
1, 1	2, 9, 16, 23, 31, 38, 45, 52, 60, 67, 74, 81
2, 1	3, 10, 17, 24, 32, 39, 46, 53, 61, 68, 75, 82
3, 1	4, 11, 18, 25, 33, 40, 47, 54, 62, 69, 76, 83
4, 1	5, 12, 19, 26, 34, 41, 48, 55, 63, 70, 77, 84
5, 1	6, 13, 20, 27, 35, 42, 49, 56, 64, 71, 78, 85
6, 1	7, 14, 21, 28, 36, 43, 50, 57, 65, 72, 79, 86
7, 1	8, 15, 22, 29, 37, 44, 51, 58, 66, 73, 80, 87

Column 1 = STS-1 POH
30 = Fixed Stuff
59 = Fixed Stuff

Figure 3-17. VT6 Locations

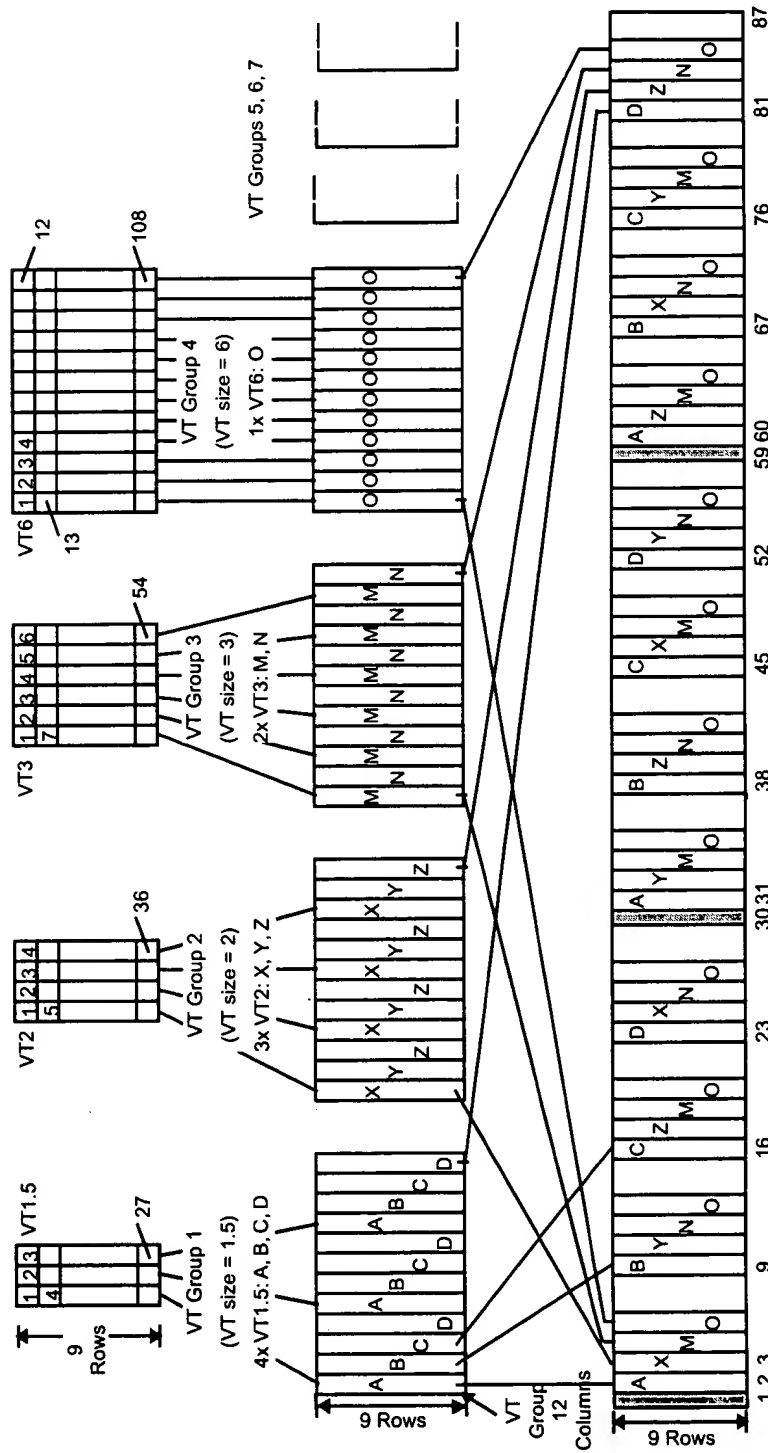


Figure 3-18. Example of VT Structured STS-1 SPE

<u>Label</u>	<u>VT Group #. VT #</u>
A	1, 1
B	1, 2
C	1, 3
D	1, 4
X	2, 1
Y	2, 2
Z	2, 3
M	3, 1
N	3, 2
O	4, 1

Figure 3-19. Correspondence Between Labels and Numbers for the Example
in Figure 3-18

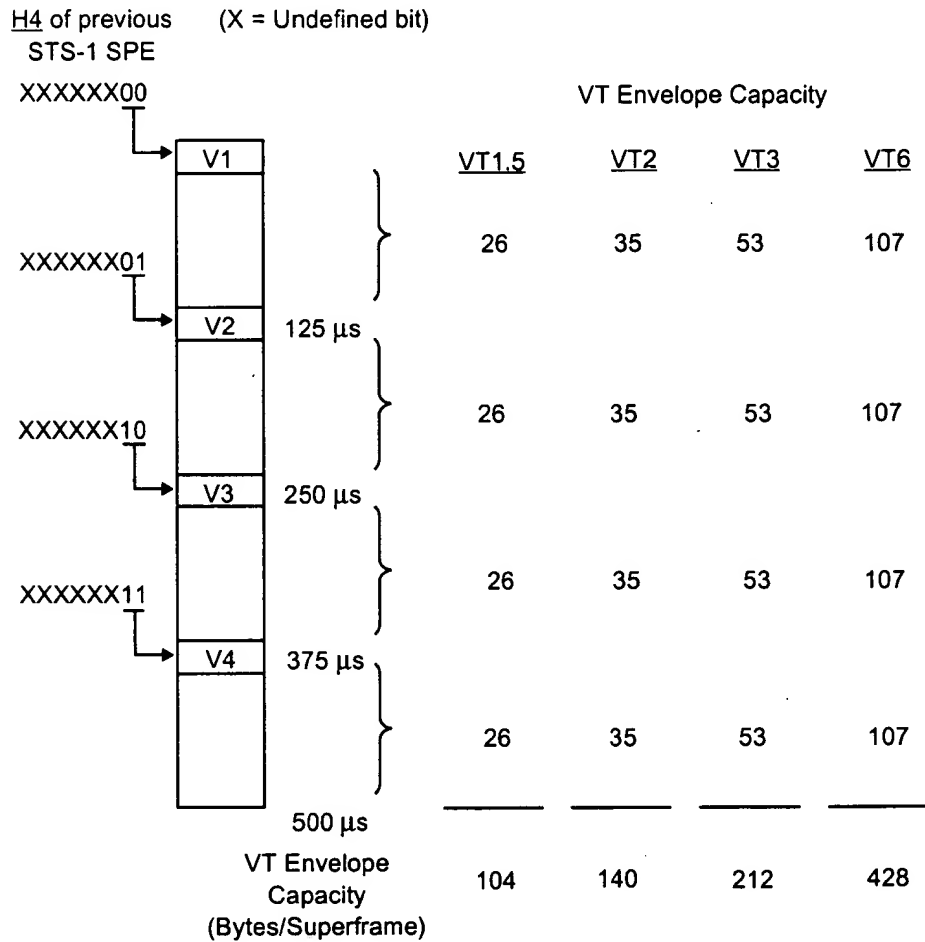


Figure 3-20. VT Superframe and Envelope Capacity

VT Superframe

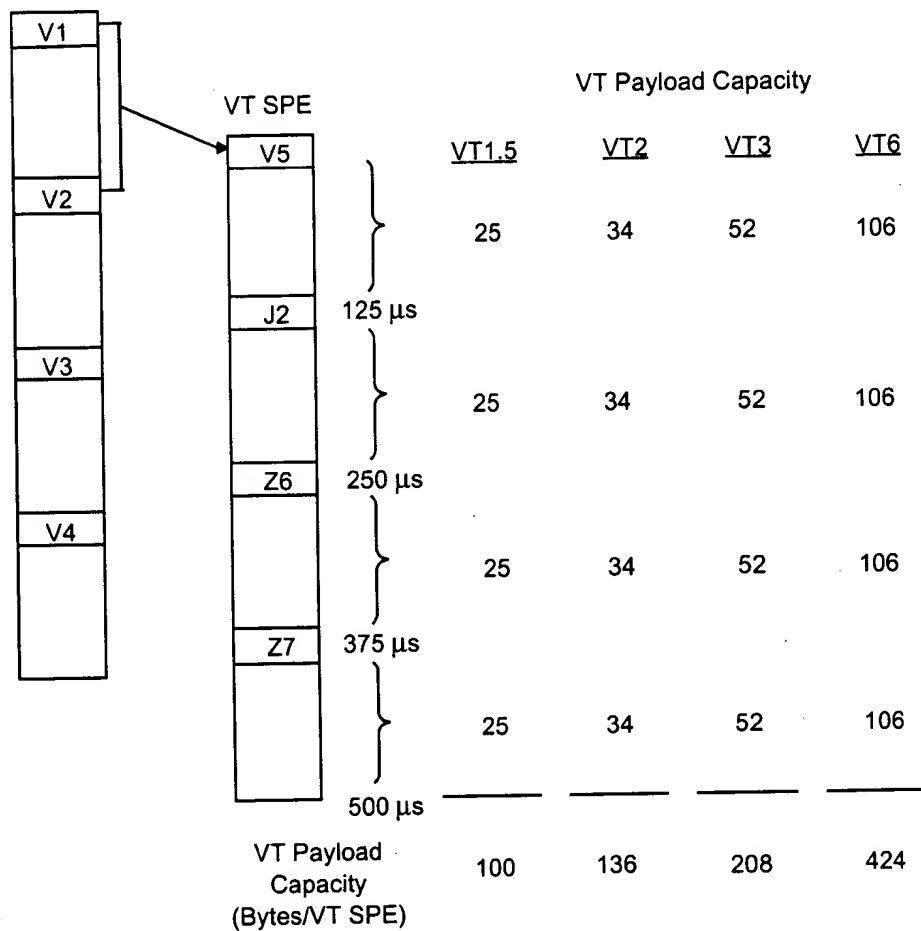


Figure 3-21. VT SPE and Payload Capacity

3.3 Layered Overhead and Transport Functions

The overhead and transport functions are broken into layers: Physical, Section, Line, and Path (see Figures 2-1 and 2-2).³ The layers have a hierarchical relationship and are considered from the top down to provide a general introduction to the individual layers and their functionalities.

Each layer requires the services of all lower-level layers to perform its functions (see Figure 3-22). For example, when two Path layer processes exchange DS3s, the Path layer maps the DS3 signal and the STS POH into an STS-1 SPE, which is then given (as an internal Path layer signal) to the Line layer. The Line layer multiplexes several SPEs from the Path layer (frame and frequency aligning each one) and adds Line overhead. Finally, the Section layer adds Section overhead and performs scrambling before transmission by the Physical layer.⁴

3.3.1 SONET Interface Layers

This section describes each layer in detail. Each description includes a broad classification of the layer, followed by a specification of the main functions it provides. Finally, examples of system hardware associated with the layer are given to clarify the role it plays. Figure 3-22 depicts the relationship of the layers to each other.

3.3.1.1 Physical Layer

The Physical layer deals with the transport of bits as optical or electrical pulses across the physical medium. No overhead is associated with the Physical layer.

The main function of this layer is conversion between internal STS-N signals and external optical or electrical SONET signals. Issues dealt with at this layer include pulse shape, power levels, and line code. As an example, electro-optical units communicate at this level.

3. Note that for payloads carried using VT-structured STS-1 SPEs, the Path layer discussed in this section actually consists of the STS-1 Path layer, and the VT Path layer, each of which performs its own functions and has its own associated overhead.

In addition, a Tandem Connection sub-layer has been defined by ANSI T1.105.05, *Synchronous Optical Network (SONET): Tandem Connection Maintenance*. When invoked, this sub-layer occurs between the Line and Path Layers and provides specific performance monitoring capabilities. This sub-layer will not be discussed further in this document, since the BCCs do not require its use.

4. Although this description (e.g., that the Line layer "adds" the Line overhead) could be interpreted to mean that the lower layer overhead bytes are not present in the signals passed from one layer to the next lower layer, that is not the intent. There are no criteria concerning the format of the internal signals used in an NE, and some (or all) of the overhead bytes may be present in the internal signals. If they are present, then it would be the function of the lower layer to overwrite those bytes as necessary to create the appropriate signal to pass to the next layer.

3.3.1.2 Section Layer

The Section layer deals with the transport of an STS-N frame across the physical medium. This layer uses the Physical layer for transport.

Functions of this layer include framing, scrambling, section error monitoring, and section level communications overhead [e.g., Local Orderwire (LOW)]. The Section overhead is interpreted and modified or created by Section Terminating Equipment (STE).

The Section and Physical layers can be used in some equipment [e.g., the STE regenerator described in TR-NWT-000917, *SONET Regenerator (SONET RGTR) Equipment Generic Criteria*] without involving the higher layers.

3.3.1.3 Line Layer

The Line layer deals with the transport of Path layer payloads across the physical medium. All lower layers exist to provide transport for this layer.

This layer provides synchronization and multiplexing functions for the Path layer. The overhead associated with these functions includes overhead for maintenance and line protection purposes and is inserted into the Line overhead channels. The Line overhead is interpreted and modified or created by Line Terminating Equipment (LTE). To access the Line overhead, the Section overhead must first be terminated. Therefore, an NE that contains Line Terminating Equipment will also contain Section Terminating Equipment.

An example of system equipment that communicates at this level is an OC-M to OC-N multiplex.

3.3.1.4 Path Layer

The Path layer deals with the transport of various payloads between SONET terminal multiplexing equipment. Examples of such payloads are DS1s and DS3s.

The Path layer maps the payloads into the format required by the Line layer. In addition, this layer communicates end-to-end via the Path Overhead (POH). The POH is interpreted and modified or created by Path Terminating Equipment (PTE). To access the Path overhead, the Section and Line overhead must first be terminated. Therefore, an NE that contains Path Terminating Equipment will also contain Section and Line Terminating Equipment.

An example of the system equipment that communicates at this level is DS3 to STS-1 mapping circuits.

3.3.1.5 Interaction of the Layers

Figure 3-22 depicts the interaction of the layers for the case of an optical interface. Each layer:

- Communicates horizontally to peer equipment in that layer
- Processes certain information and passes it vertically to the adjacent layers.

The interactions are described in terms of each level's horizontal and vertical transactions.

Figure 3-22 also shows payloads as inputs to the Path layer. This layer transmits the payloads and the POH horizontally to its peer entities. The Path layer maps the payloads and POH into SPEs that it passes vertically to the Line layer as internal Path layer signals.

The Line layer transmits the SPEs and Line overhead to its peer entities. It maps the SPEs and Line overhead into internal Line layer signals. The SPEs are synchronized and multiplexed at this time, and then the internal Line layer signal is passed to the Section layer.

The Section layer transmits STS-N signals to its peer entities. It maps the internal Line layer signals and the Section overhead into an internal STS-N signal that is handed to the Physical layer, which transmits optical or electrical pulses to its peer entities.

Access to all of the layers is not required of every SONET NE. For example, an STE regenerator would use only the first two layers (Physical and Section). Similarly, an NE that merely routes SPEs and does not accept any new inputs from the Path layer uses only the first three layers (Physical, Section, and Line). Note however, that NEs may monitor (or in some cases, may be required to monitor) the overhead of layers that they do not terminate.

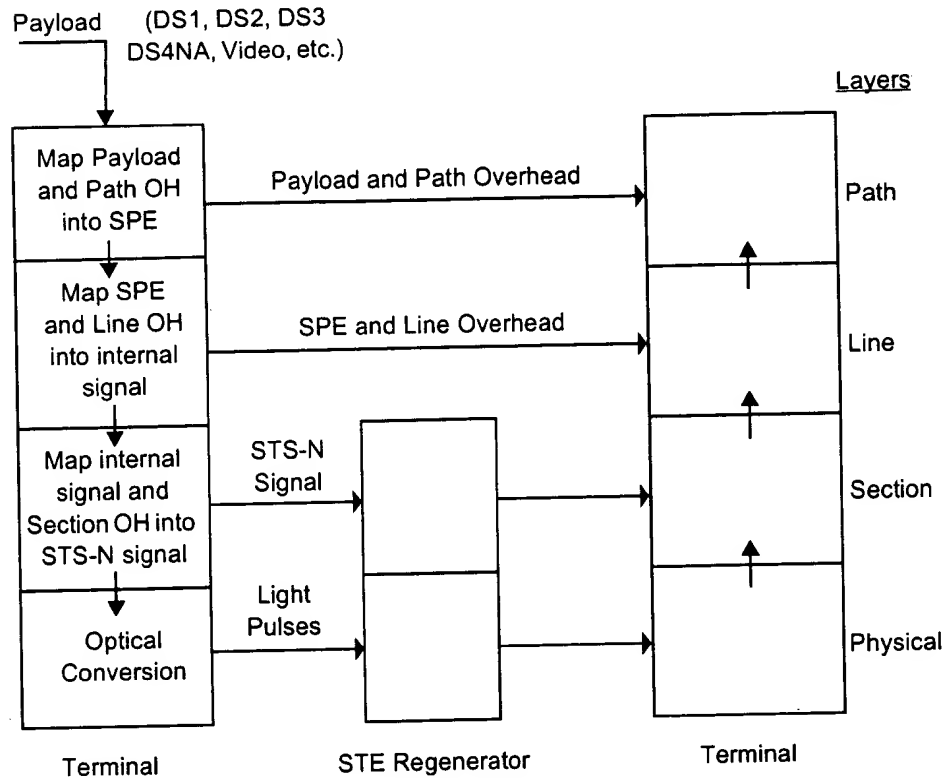


Figure 3-22. Optical Interface Layers

3.3.2 STS-1 Overhead Descriptions

Figure 3-5 illustrated the location of the overhead bytes in the STS-1 frame. The functions assigned to the Section and Line layers have been combined into a structure of 27 bytes called the Transport overhead, which occupies the first 3 columns of the frame. The functions of the STS Path layer have been assigned nine bytes in the first column of the STS SPE.

This section defines each of the overhead bytes. Each byte is assigned to a layer and a position in the overhead columns as shown in Figure 3-23. The overhead associated with a given layer is modified or created by the equipment terminating that layer before insertion on the outgoing signal.

In addition, two types of interface signals are defined. These are line-side signals, which have full Transport overhead functionality, and drop-side signals, which have reduced

Transport overhead functionality. Line-side signals are suitable for inter-office connections, although they could also be used for intra-office connections. Drop-side signals are suitable for intra-office connections. The overhead functionality criteria applicable to line-side and drop-side signals are summarized in Table 5-2.

The descriptions start with the Section layer, because there is no overhead associated with the Physical layer.

Transport Overhead				Path Overhead
Section Overhead	Framing	Framing	Trace/Growth (STS-ID) J0/Z0 ^a	Trace
	A1	A2	User	J1
	BIP-8	Orderwire	User	BIP-8
	B1/Undefined ^a	E1/Undefined ^a	F1/Undefined ^a	B3
	Data Com	Data Com	Data Com	Signal Label
Line Overhead	D1/Undefined ^a	D2/Undefined ^a	D3/Undefined ^a	C2
	Pointer	Pointer	Pointer Action	Path Status
	H1	H2	H3	G1
	BIP-8	APS	APS	User Channel
	B2	K1/Undefined ^a	K2/Undefined ^a	F2
	Data Com	Data Com	Data Com	Indicator
	D4/Undefined ^a	D5/Undefined ^a	D6/Undefined ^a	H4
	Data Com	Data Com	Data Com	Growth
	D7/Undefined ^a	D8/Undefined ^a	D9/Undefined ^a	Z3
	Data Com	Data Com	Data Com	Growth
	D10/Undefined ^a	D11/Undefined ^a	D12/Undefined ^a	Z4
	Sync Status/ Growth S1/Z1 ^a	REI-L ^b /Growth M0 or M1/Z2 ^a	Orderwire E2/Undefined ^a	Tandem Connection Z5

- a. For entries of the form "X/Y", the first label shown is applicable for one STS-1 in an STS-N electrical or OC-N signal, and the second label is applicable for the remaining STS-1s.
- b. REI-L (Line Remote Error Indication) was previously referred to as Line FEBE.

Figure 3-23. Transport and Path Overhead Byte Designations

3.3.2.1 Section Overhead

This section defines each of the Section overhead bytes.

Framing (A1 and A2) – Two bytes are allocated in each STS-1 for framing.

- R3-13** [16] The A1 byte shall be set to '11110110' and the A2 byte shall be set to '00101000' in all STS-1s within an STS-N.

Section 5.5 contains the framing criteria for SONET NEs.

Section Trace (J0)/Section Growth (Z0) – The byte in each of the N STS-1s in an STS-N that was formerly defined as the STS-1 ID (C1) byte has been redefined either as the Section Trace byte (in the first STS-1 of the STS-N), or as a Section Growth byte (in the second through Nth STS-1s). Detailed criteria concerning the use of these bytes for their new functions are for further study. Until those details are determined, the following criteria apply. They will be modified as necessary when the details of the Section Trace feature and uses for the Section Growth bytes are defined.

- O3-14** [17] STE that supports line-side signals should have the capability to access the J0 byte, which is located in the first STS-1 of an STS-N.

The ability to access the J0 byte is not required for STE that only supports drop-side signals.

- R3-15** [18] Unless it is being used for a defined purpose (e.g., to carry a Section Trace message once the details of that feature are defined) each J0 and Z0 byte shall be set to a binary number corresponding to its order of appearance in the STS-N frame (i.e., the J0 byte shall be set to 00000001, the first Z0 byte shall be set to 00000010, the second Z0 byte to 00000011, etc.).

The preceding requirement is applicable for both line-side and drop-side signals. Also, since no standard use has been defined for the J0 and Z0 bytes (or for the former C1 bytes) received by an NE, these bytes are considered unused at the receiving STE (see Section 3.2).

Section BIP-8 (B1) – The B1 byte is located in the first STS-1 of an STS-N, and is defined for a Section error monitoring function in line-side signals. The value contained in the B1 byte in a drop-side signal is undefined, so the criteria in Section 3.2 are applicable.⁵ The corresponding byte locations in the second through Nth STS-1s of both line-side and drop-side signals are also currently undefined.

5. It is also acceptable for the B1 byte in a drop-side signal to carry the BIP-8 code as described for line-side signals. However, STE cannot assume the BIP-8 code will be present in the received drop-side signal, and therefore it must be capable of ignoring the value in that byte.

- R3-16** [19] The B1 byte in a line-side signal shall carry a BIP-8 code, using even parity. The Section BIP-8 shall be calculated over all bits of the previous STS-N frame after scrambling and placed in the B1 byte of the current STS-N frame before scrambling.

Orderwire (E1) – The E1 byte is located in the first STS-1 of an STS-N, and is used for an LOW channel. The corresponding byte locations in the second through Nth STS-1s are currently undefined. The LOW channel is used for voice communication between regenerators, hubs, and remote terminal locations. Section 5.2.2 contains the Orderwire criteria.

Section User Channel (F1) – The F1 byte is located in the first STS-1 of an STS-N, and is available for use by the network provider. The corresponding byte locations in the second through Nth STS-1s are currently undefined. Section 5.2.3 contains the Section User Channel criteria.

Section Data Communication Channel (D1, D2, and D3) – The D1, D2, and D3 bytes are located in the first STS-1 of an STS-N, and are used for Section data communications. The corresponding byte locations in the second through Nth STS-1s are currently undefined.

These three bytes are considered as one 192-kb/s, message-based channel for alarms, maintenance, control, monitoring, administering, and other communication needs between STE. This channel is used for internally generated, externally generated, and supplier-specific messages. Section 8 contains the data communications channel criteria.

3.3.2.2 Line Overhead

This section defines each of the Line overhead bytes.

STS Payload Pointer (H1 and H2) – Two bytes are allocated to a pointer that indicates the offset in bytes between the pointer and the first byte of the STS SPE. The pointer bytes are used in all STS-1s within an STS-N to align the STS-1 Transport Overheads in the STS-N, and to perform frequency justification (see Section 3.5).

These bytes are also used to indicate concatenation, and to detect STS Path Alarm Indication Signals (AIS-P).

Pointer Action Byte (H3) – The pointer action byte is allocated for SPE frequency justification purposes. The H3 byte is used in all STS-1s within an STS-N to carry the extra SPE byte in the event of a negative pointer adjustment. The value contained in this byte when it is not used to carry the SPE byte is undefined.

Line BIP-8 (B2) – One byte is allocated in each STS-1 for a Line error monitoring function. The N Line BIP-8 bytes in an STS-N electrical or OC-N signal are intended to form a single error monitoring facility capable of measuring bit error ratios up to 10^{-3} , independent of the value of N.

- R3-17** [20] The B2 byte shall be provided in all STS-1s within an STS-N to carry a Line BIP-8 code, using even parity. The Line BIP-8 shall be calculated over all bits of the Line Overhead and the Envelope Capacity of the previous STS-1 frame before scrambling, and placed in the B2 byte of the current STS-1 frame before scrambling.

APS Channel (K1 and K2) – The K1 and K2 bytes are located in the first STS-1 of an STS-N, and are used on the protection line for Automatic Protection Switching (APS) signaling between LTE that uses line level protection switching (e.g., in systems using linear APS, or in bidirectional line switched rings). The corresponding byte locations in the second through Nth STS-1s are currently undefined.

The K2 byte is also used to detect Line AIS (AIS-L) and Line Remote Defect Indication (RDI-L) signals (see Sections 6.2.1.2.1 and 6.2.1.3.1).

Line Data Communication Channel (D4 through D12) – The D4 through D12 bytes are located in the first STS-1 of an STS-N, and are used for Line data communication. The corresponding byte locations in the second through Nth STS-1s are currently undefined.

These nine bytes are considered as one 576-kb/s, message-based channel for alarms, maintenance, control, monitoring, administering, and other communication needs. This channel is available for internally generated, externally generated, and supplier-specific messages. Section 8 contains the data communications channel criteria.

Synchronization Status (S1) – The S1 byte is located in the first STS-1 of an STS-N, and bits 5 through 8 of that byte are allocated to convey the synchronization status of the NE. Section 5.4.2 contains the synchronization status message criteria. Bits 1 through 4 of the S1 byte are currently undefined.

Growth (Z1) – The Z1 bytes are located in the second through Nth STS-1s of an STS-N ($3 \leq N \leq 48$), and are allocated for future growth. The use of these bytes is currently undefined. Note that an OC-1 or STS-1 electrical signal does not contain a Z1 byte.

STS-1 REI-L (M0) – The M0 byte is defined only for the STS-1 in an OC-1 or STS-1 electrical signal. Bits 5 through 8 of the M0 byte are allocated for a Line Remote Error Indication function (REI-L, formerly referred to as Line FEBE), which conveys the error count detected by LTE (using the Line BIP-8 code) back to its peer LTE. Bits 1 through 4 of the M0 byte are currently undefined.

- R3-18** [21] LTE terminating an OC-1 or STS-1 electrical signal shall set bits 5 through 8 of the M0 byte to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected based on the Line BIP-8 (B2) byte. The error count shall be a binary number from zero (i.e., 0000) to 8 (i.e., 1000). The remaining seven values represented by the four REI-L bits (i.e., 1001 through 1111) shall not be transmitted, and shall be interpreted by receiving LTE as zero errors.

STS-N REI-L (M1) – The M1 byte is located in the third STS-1⁶ (in order of appearance in the byte-interleaved STS-N electrical or OC-N signal) in an STS-N ($N \geq 3$), and is used for an REI-L function.

- R3-19** [22] LTE terminating an OC-N or STS-N electrical signal ($N \geq 3$) shall set the M1 byte to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected using the Line BIP-8 (B2) bytes. For values of N below 48, the error count shall be a binary number from zero to 8N. The remaining possible values [i.e., $255 - (8 \times N)$] represented by the eight REI-L bits shall not be transmitted, and shall be interpreted by the receiving LTE as zero errors. For N equal to 48, the count shall be truncated at 255.

Growth (Z2) – The Z2 bytes are located in the first and second STS-1s of an STS-3, and the first, second, and fourth through Nth STS-1s of an STS-N ($12 \leq N \leq 48$). These bytes are allocated for future growth, and their use is currently undefined. Note that an OC-1 or STS-1 electrical signal does not contain a Z2 byte.

Orderwire (E2) – The E2 byte is located in the first STS-1 of an STS-N, and is used for an Express Orderwire (EOW) channel between Line entities. The corresponding byte locations in the second through Nth STS-1s are currently undefined. Section 5.2.2 contains the Orderwire criteria.

3.3.2.3 STS Path Overhead (STS POH)

STS POH is assigned to each STS SPE for functions necessary in transporting its payload. The STS POH supports the following classes of functions:

- A. Payload independent functions with standard format and coding – all payloads require these functions.
- B. Mapping dependent functions with standard format and coding that are specific to the type of payload – these functions are needed for one or more types of payload, but not all payloads.

6. The “third STS-1” is defined as the third STS-1 in order of their appearance in the byte-interleaved STS-N electrical or OC-N signal. Using the two-level STS numbering scheme discussed in Figures 5-1 and 5-2, the third STS-1 in an OC-12 or higher rate signal would be labeled “3,1”.

It is important to recognize that the numbering of the transport overhead bytes as they appear in an STS-N electrical or OC-N signal can be separated from the numbering of the STS-1 and STS-M inputs to the byte-interleaver (see Section 5.1.1). When the NE adds (or overwrites) the Line Overhead to the output of the byte-interleaver, it does so independently of the particular mix of input STS-1s and STS-Ms to the byte-interleaver. For example, if the first input to the byte-interleaver shown in Figure 5-2 was “STS-12 Number 1,1” (instead of “STS-1 Number 1,1”) then there would be no inputs shown with numbers of “1,2”, “1,3”, “2,1”, ..., “4,3” (including “3,1”). However, the M1 byte would still appear in a particular transport overhead byte position, and that byte position can be said to be contained in STS-1 “3,1”.

- C. Application-specific functions – appropriate GRs, TRs, TAs, or standards documents (e.g., ANSI T1.105.05 for applications using the Tandem Connection sub-layer) specify the format and coding for these functions, which may share the same overhead capacity.

STS POH capacity not yet assigned to Class A, B, or C functions may be defined in the future for supporting any of those classes, with Class A having priority. Also, this classification scheme does not preclude the allocation of other overhead bits or bytes within the STS Payload Capacity for specific mappings (e.g., the stuff control bits for the asynchronous DS3 mapping).

Class A Functions

STS Path Trace (J1) – This byte is used to transport a repetitive 64-byte message so that receiving STS PTE can verify its continued connection to the intended transmitting STS PTE.

Sections 6.2.1.1.9 and 6.2.3.2.3.A contain the criteria related to loading and detecting the STS Path Trace message. In general, an 8-bit ASCII CLLI™ code, padded with ASCII NULL characters and terminated with CR and LF characters (for a total of 64 bytes) would be a suitable message.

- R3-20** [23] If no message has been loaded by the user for transmission in the J1 byte, then that byte shall be set to all-zeros (i.e., to ASCII NULL characters).

STS Path BIP-8 (B3) – The B3 byte is allocated for an STS Path error monitoring function.

- R3-21** [24] The B3 byte shall carry a BIP-8 code, using even parity. The STS Path BIP-8 shall be calculated over all bits (783 bytes for an STS-1 SPE or N×783 bytes for an STS-Nc SPE, regardless of any pointer adjustments) of the previous STS SPE before scrambling, and placed in the B3 byte of the current STS SPE before scrambling.

STS Path Signal Label (C2) – The C2 byte is allocated to indicate the content of the STS SPE, including the status of the mapped payloads. Of the 256 possible binary values [00 to FF (hex)], only the codes defined in Tables 3-2 and 3-3 have been assigned for the mappings defined in Section 3.4. The codes in Table 3-2 are assigned for use by all STS PTE, and the code generated (under normal conditions) by a particular STS PTE depends on its provisioned (or only supported) functionality. The additional codes in Table 3-3 are assigned for use by STS PTE that support the STS Payload Defect Indication (PDI-P) feature, and are generated automatically based on the status of the mapped payloads. The remaining codes are reserved to be assigned as needed for future STS payload-specific mappings.

PDI-P is an application-specific feature that uses the STS Path Signal Label to indicate to downstream equipment that there is a defect in one or more directly mapped, embedded payloads in the STS SPE. For VT-structured STS-1 SPEs, 28 codes are defined to indicate that 1 through 28 of the embedded VTxs have defects. In this scheme a VT1.5 payload defect is assigned the same priority as a VT2, VT3, or VT6 payload defect. For example, an STS-1 SPE with two VT6 payload defects would be assigned the code 'E2,' which is the same code as would be used for an STS-1 SPE with two VT1.5 payload defects. For non-VT-structured STS-1 or STS-Nc SPEs, one code is defined to indicate that the mapped payload has a defect. See Section 6.2.1.4.1 for additional criteria related to PDI-P.

R3-22 [25] For an STS path connection that is equipped and provisioned, a valid non-zero STS Path Signal Label shall be generated by the STS PTE. If the content of the STS SPE is one of the specific possibilities listed in Table 3-2, then the corresponding code from Table 3-2 (or if PDI-P is supported and one or more Payload Defects are present, Table 3-3) shall be used. If the content is not specifically listed, then the code for "Equipped - Nonspecific Payload" shall be used.

R3-23 [26] For STS path connections that are not equipped, or that are equipped but not provisioned, the NE (e.g., the NE's LTE) shall generate all-zeros STS SPEs with "valid" STS Payload Pointers.

In the above requirements, "provisioned" means that the STS PTE has been configured for a mapping (or only supports one mapping), and has been assigned a timeslot (or is hardwired to a specific timeslot) in a SONET signal. Unless both of these conditions are met, the path connection is not considered provisioned. In addition, a supplier may choose not to consider a path connection to be provisioned unless the STS PTE has been assigned one or more signals (e.g., one or more VTs) to map or multiplex into the STS SPE that it is originating.

Note that Issue 1 of this document required that a path connection not be considered provisioned unless the STS PTE had been assigned a signal (rather than leaving it as a supplier option). That meant (for example) that STS PTE terminating a VT-structured STS-1 SPE was not considered provisioned unless it had been assigned one or more specific VTs to place into the STS-1 SPE that it was originating. However, in some applications it may be useful to confirm STS PTE to STS PTE connectivity before they have been assigned specific signals to map or multiplex. NEs that provide that capability [e.g., by sending non-zero STS signal labels or STS Path Trace messages (along with the appropriate STS Path BIP-8) before being assigned a signal] should not be considered nonconforming, and therefore the definition of provisioned was revised.

In addition, note that the "valid" pointers referred to in **R3-23** [26] may contain any value that can be interpreted correctly by downstream STS pointer processors without causing STS Loss of Pointer (see Section 6.2.1.1.3), and that an all-zeros STS SPE results in a valid STS Path BIP-8 code and a signal label of '00' (hex), indicating "Unequipped".

Table 3-2. STS Path Signal Label Assignments

Code (Hex)	Content of the STS SPE
00	Unequipped
01	Equipped – Nonspecific Payload
02	VT-Structured STS-1 SPE ¹
03	Locked VT Mode ¹
04	Asynchronous Mapping for DS3
12	Asynchronous Mapping for DS4NA
13	Mapping for ATM
14	Mapping for DQDB
15	Asynchronous Mapping for FDDI
16	HDL-Over-SONET Mapping
FE	O.181 Test Signal (TSS1 to TSS3) Mapping ²

Note:

1. In previous SONET standards and criteria documents, two modes were defined for VT-structured STS-1 SPEs. These were the floating mode, which was assigned the code '02,' and the locked mode, which was assigned '03.' The Locked VT Mode has since been removed from the SONET standards and criteria. However, for backward compatibility between future mappings and equipment that supports the Locked VT Mode, the signal label that was assigned to that mode remains defined.
2. This code/mapping assignment is specified in ITU-T Recommendation G.707, *Network node interface for the Synchronous Digital Hierarchy (SDH)*, for use in SDH networks.

Table 3-3. STS Path Signal Label Assignments for Signals with Payload Defects

Code (Hex)	Content of the STS SPE	Code (Hex)	Content of the STS SPE
E1	VT-structured STS-1 SPE with 1 VTx Payload Defect (STS-1 w/1 VTx PD)	F0	STS-1 w/16 VTx PDs
		F1	STS-1 w/17 VTx PDs
E2	STS-1 w/2 VTx PDs	F2	STS-1 w/18 VTx PDs
E3	STS-1 w/3 VTx PDs	F3	STS-1 w/19 VTx PDs
E4	STS-1 w/4 VTx PDs	F4	STS-1 w/20 VTx PDs
E5	STS-1 w/5 VTx PDs	F5	STS-1 w/21 VTx PDs
E6	STS-1 w/6 VTx PDs	F6	STS-1 w/22 VTx PDs
E7	STS-1 w/7 VTx PDs	F7	STS-1 w/23 VTx PDs
E8	STS-1 w/8 VTx PDs	F8	STS-1 w/24 VTx PDs
E9	STS-1 w/9 VTx PDs	F9	STS-1 w/25 VTx PDs
EA	STS-1 w/10 VTx PDs	FA	STS-1 w/26 VTx PDs
EB	STS-1 w/11 VTx PDs	FB	STS-1 w/27 VTx PDs
EC	STS-1 w/12 VTx PDs	FC	VT-structured STS-1 SPE with 28 VT1.5 Payload Defects, or a non-VT- structured STS-1 or STS-Nc SPE with a Payload Defect
ED	STS-1 w/13 VTx PDs		
EE	STS-1 w/14 VTx PDs		
EF	STS-1 w/15 VTx PDs		

Path Status (G1) – The G1 byte is allocated to convey the Path terminating status and performance back to the originating STS PTE. This feature permits the status and performance of the complete duplex path to be monitored at either end, or at any point along that path. Figure 3-24 illustrates the bit assignments in the G1 byte. Bits 1 through 4 are allocated for an STS Path REI function (REI-P, formerly referred to as STS Path FEBE).

- R3-24** [27] STS PTE shall set bits 1 through 4 of the G1 byte to indicate (to the upstream STS PTE) the count of interleaved-bit block errors that it has detected based on the STS Path BIP-8 byte (B3). The error count shall be a binary number from zero (i.e., 0000) to 8 (i.e., 1000). The remaining seven values represented by the four REI-P bits (i.e., 1001 through 1111) shall not be transmitted, and shall be interpreted by receiving STS PTE as zero errors.

Bits 5, 6 and 7 of the G1 byte are allocated for an STS Path RDI (RDI-P) signal.
Section 6.2.1.3.2 contains the criteria related to the generation and detection of RDI-P.
Bit 8 of the G1 byte is currently undefined.

REI-P				RDI-P			Und.
1	2	3	4	5	6	7	8

REI-P Coding:

0	0	0	0	0	Errors
0	0	0	1	1	Error
0	0	1	0	2	Errors
0	0	1	1	3	Errors
0	1	0	0	4	Errors
0	1	0	1	5	Errors
0	1	1	0	6	Errors
0	1	1	1	7	Errors
1	0	0	0	8	Errors
1	0	0	1	} 0 Errors	
1	0	1	0		
1	0	1	1		
1	1	0	0		
1	1	0	1		
1	1	1	0		
1	1	1	1		

Figure 3-24. STS Path Status Byte (G1)

Class B Functions

Indicator (H4) – The H4 byte is allocated for use as a mapping-specific indicator byte. Currently, it is used only in VT-structured STS-1 SPEs and the DQDB mapping. For VT-structured STS-1 SPEs, the byte is used as a multiframe indicator. See Section 3.4.1 for the applicable criteria. In the DQDB mapping it is used to carry the DQDB Link Status Signal, and to indicate the offset to the boundary of the next DQDB slot. See Section 3.4.3.1.4 for details.

For mappings where a standard use of this byte has not been defined, it is considered undefined.

Class C Functions

Path User Channel (F2) – The F2 byte is allocated for user communication purposes between STS Path terminating NEs. For applications where this byte is not used, it is considered undefined. Section 5.2.3 contains the criteria for the Path User Channel.

The F2 byte is also used in the DQDB mapping to carry DQDB Layer Management information (a Class B function). See Section 3.4.3.1.4 for details.

Tandem Connection (Z5) – The Z5 byte is allocated for Tandem Connection Maintenance and the Path Data Channel. Refer to ANSI T1.105 and ANSI T1.105.05 for details.

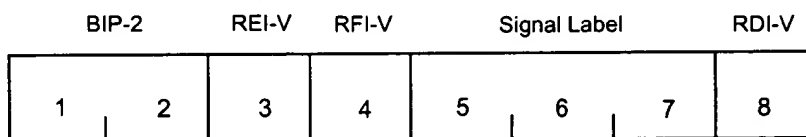
Other

STS Path Growth (Z3, Z4) – The Z3 and Z4 bytes are allocated for future growth, and the use of these bytes is currently undefined for most mappings and applications. In the DQDB mapping the Z3 byte is used to carry DQDB Layer Management information (a Class B function). See Section 3.4.3.1.4 for details.

3.3.3 VT Path Overhead

Four bytes (V5, J2, Z6, and Z7) are allocated for VT POH. The first byte of a VT SPE (i.e., the byte in the location pointed to by the VT Payload Pointer) is the V5 byte, while the J2, Z6, and Z7 bytes occupy the corresponding locations in the subsequent 125- μ s frames of the VT Superframe (as shown in Figure 3-21).

V5 – The V5 byte provides the same functions for VT paths that the B3, C2, and G1 bytes provide for STS paths; namely error checking, signal label, and path status. The bit assignments for the V5 byte are illustrated in Figure 3-25.



REI-V Coding

0	0 Errors
1	1 or 2 Errors

Figure 3-25. VT Path Overhead Byte (V5)

Bits 1 and 2 of the V5 byte are allocated for error performance monitoring. A BIP-2 scheme is defined as follows:

- R3-25** [28] Bit 1 of the V5 byte shall be set so that the parity of all of the odd-numbered bits (i.e., bits 1, 3, 5, and 7) in all bytes in the previous VT SPE is even. Bit 2 shall be set so that the parity of all of the even-numbered bits (2, 4, 6, and 8) in all bytes in the previous VT SPE is even.

Bit 3 of the V5 byte is allocated for a VT Path REI function (REI-V, formerly referred to as VT Path FEBE) to convey the VT Path terminating performance back to an originating VT PTE.

- R3-26** [29] VT PTE shall set bit 3 of the V5 byte to '1' if one or more errors were detected using the BIP-2. It shall set bit 3 to '0' if zero errors were detected.

Bit 4 of the V5 byte is allocated for a VT Path Remote Failure Indication (RFI-V) in the byte-synchronous DS1 mapping. See Section 6.2.1.3.3 for the criteria related to the generation and detection of RFI-V signals.

Bits 5 through 7 of the V5 byte are allocated for a VT Path Signal Label to indicate the content of the VT SPE. Of the 8 possible binary values ('000' to '111'), only the codes defined in Table 3-4 for each VT size have been assigned for the mappings described in Section 3.4.1. The remaining codes are reserved to be assigned as required for future VT payload mappings.

- R3-27** [30v2] For a VT path connection that is equipped and provisioned, a valid, non-zero VT Signal Label shall be generated by the VT PTE. If the mapping contained in the VT SPE is one of the specific possibilities listed in Table 3-4, then the corresponding code from the table shall be used. If the mapping is not specifically listed, then the code for "Equipped - Nonspecific Payload" shall be used.

- R3-28** [31] For VT path connections that are not equipped, or that are equipped but not provisioned, the NE (e.g., the NE's STS PTE) shall generate all-zeros VT SPEs with "valid" VT Payload Pointers.

In the above requirements, "provisioned" means that the VT PTE has been configured for a mapping (or only supports one mapping), and has been assigned a timeslot (or is hardwired to a specific timeslot) in an STS-1 SPE. Unless both of these conditions are met, the path connection is not considered provisioned. In addition, a supplier may choose not to consider a path connection to be provisioned unless the VT PTE has been assigned a signal (e.g., via a cross-connection) to map or multiplex into the VT SPE that it is originating.

In addition, note that the "valid" pointers referred to in R3-28 [31] may contain any value that can be interpreted correctly by downstream VT pointer processors without causing VT Loss of Pointer (see Section 6.2.1.1.3), and that an all-zeros VT SPE results in a valid VT Path BIP-2 code and a signal label of '000', indicating "Unequipped".

Table 3-4. VT Signal Label Assignments

Code	VT1.5	VT2 ^a	VT3	VT6
000	Unequipped			
001	Equipped – Nonspecific Payload			
010	Asynchronous Mapping for DS1	Asynchronous Mapping for 2.048 Mb/s	Asynchronous Mapping for DS1C	Asynchronous Mapping for DS2
011	Bit-synchronous Mapping for DS1 ^b	Bit-synchronous Mapping for 2.048 Mb/s		
100	Byte-synchronous Mapping for DS1	Byte-synchronous Mapping for 2.048 Mb/s		

Notes:

- Although the VT2 mappings for 2.048 Mb/s signals are not included in this document, the signal label codes for those mappings have been assigned in ANSI T1.105.
- The DS1 bit-synchronous mapping has been removed from the SONET standards and criteria. However, for backward compatibility between future mappings and equipment that supports the DS1 bit-synchronous mapping, the signal label that was assigned to that mapping will remain defined.

Bit 8 of the V5 byte (along with bits 5 through 7 of the Z7 byte, see below) is allocated for a VT Path Remote Defect Indication (RDI-V) signal. Section 6.2.1.3.3 contains the criteria related to the generation and detection of RDI-V.

VT Path Trace (J2) – The J2 byte is allocated for a VT Path Trace function; however, the details of that function are for further study. Therefore, the byte is currently considered undefined.

VT Path Growth (Z6) – The Z6 byte is allocated for future growth.

VT Path Growth (Z7) – Bits 1, 2, 3, 4, and 8 of the Z7 byte are allocated for future growth. Bits 5 through 7 of the Z7 byte (along with bit 8 of the V5 byte, see above) are allocated for an RDI-V signal. Section 6.2.1.3.3 contains the criteria related to the generation and detection of RDI-V.

3.4 Payload Mapping

This section describes the mapping of various payloads into SPEs. This includes sub-STs-1 payloads, STs-1 payloads, and Super Rate (i.e., STs-Nc) payloads.

3.4.1 Sub-STs-1 Mappings

Payloads below the DS3 rate are transported in a VT structure. This section describes the payload mappings that use the VT-structured STs-1 SPE.

- R3-29** [32] The H4 byte shall be used to indicate phase of the V1 through V4 bytes in the 500- μ s (i.e., 4-frame) VT Superframe. The allocation of the bits in the H4 byte, and the correspondence of the H4 code with the V1 through V4 bytes shall be as shown in Figure 3-26.

When the H4 byte is used to indicate the phase of the V1 through V4 bytes, bits 1 through 6 are considered to be undefined. Therefore the criteria in Section 3.2 are applicable, including the objective that undefined bits and bytes be set to zero. However, ANSI T1.105 and previous Bellcore criteria documents required H4 bits 1 through 6 to be set to non-zero values. To accommodate equipment built to the standard or to the previous requirements, an NE that does either of the following will also be considered to have met the objective (for the undefined bits in H4):

- Sets bits 1 through 6 to all-ones
- Sets bits 1 through 6 to the 48-frame sequence described below:
 - All of the bits (including bits 7 and 8) are zero in frame 0
 - Bits 1 and 2 count from '00' through '11' over 24 frames (i.e., '00' appears in 6 consecutive frames, then '01' for 6 frames, etc.) and then repeat
 - Bits 3 and 4 count from '00' through '10' over 6 frames, and then repeat
 - Bits 5 and 6 count from '00' through '11' over 16 frames, and then repeat.

H4 Byte bits								V1 through V4 byte in the next frame	Time
1	2	3	4	5	6	7	8		
X	X	X	X	X	X	0	0	V1	0
X	X	X	X	X	X	0	1	V2	
X	X	X	X	X	X	1	0	V3	
X	X	X	X	X	X	1	1	V4	

500- μ s VT Superframe

Figure 3-26. H4 Byte Coding Sequence for VT-Structured STs-1 SPEs

Although there are no specific criteria related to the alignment algorithm that STS PTE uses in determining the phase of the received H4 byte sequence, it is important that the algorithm be tolerant of single bit errors in the H4 sequence, and that it gain alignment quickly after the H4 byte is located. It is also important that the algorithm quickly recognize and respond to actual changes in the phase of the H4 bytes. If STS PTE cannot determine the phase of the H4 sequence, then it also cannot determine the phase of the V1 through V4 bytes. This is likely to result in a disruption of traffic and the detection of multiple LOP-V defects (see Section 6.2.1.1.3).

Asynchronous mappings are defined for clear-channel transport of signals that meet the DSX requirements in GR-499-CORE. At asynchronous interfaces, frame acquisition and generation are not required for mapping purposes, although frame acquisition may be required for monitoring purposes (for example, see Section 6.2.2.9). Byte-synchronous mappings, which are currently defined (in this document) only for DS1 signals, allow direct identification and access to the DS0 channels that are carried. At byte-synchronous interfaces, frame acquisition and generation capabilities are required.

3.4.1.1 Byte-Synchronous Mapping for DS1

A byte-synchronous mapping of a DS1 into the payload capacity of a VT1.5 SPE is defined to allow downstream SONET NEs direct identification and access to the 24 DS0 channels that are carried. In some applications, a DS1 signal that appears at a DS1 interface is byte-synchronously mapped into the VT1.5, while in other applications the mapping is used to transport 24 DS0 channels without a DS1 interface [e.g., in an Integrated Digital Loop Carrier (IDLC) application with SONET-based Remote Digital Terminals (RDTs) and Integrated Digital Terminals (IDTs)]. In addition, some NEs may provide DS0 rearrangement (i.e., cross-connect and grooming) capabilities.

Figure 3-27 shows the byte-synchronous mapping for a DS1 (or 24 DS0s) into a VT1.5. The S₁, S₂, S₃, and S₄ bits are allocated to carry signaling for the 24 DS0 channels, the F-bit is allocated to carry the DS1 frame bit, and the P-bits are allocated for indicating the phase of the signaling and the frame bits on a per-VT basis.

R3-30 [33] If the byte-synchronous DS1 mapping is provided, it shall be as shown in Figure 3-27.

In some applications the S-bits or the F-bit are not used to carry signaling or framing (see Section 3.4.1.1.1). In those cases they are considered undefined and the criteria in Section 3.2 are applicable. Similarly, if the P-bits are not needed to indicate the phase of the S-bits or the F-bits in a particular application, then they are also considered undefined.⁷

7. It is acceptable for the VT PTE to set the P-bits to either all-zeros or the 24-frame sequence in Figure 3-28 when they are undefined.

- R3-31** [34] If the P-bits are being used to indicate the phase of the S-bits or the F-bits, then they shall set to the 24-frame sequence shown in Figure 3-28.

Although there are no specific criteria related to the alignment algorithm that VT PTE uses in determining the phase of the received P-bit sequence (and therefore the phase of the S-bits or the F-bits), it is important that the algorithm be tolerant of single bit errors in the P-bit sequence, and that it gain alignment quickly after the P-bits are located. It is also important that the algorithm quickly recognize and respond to actual changes in the phase of the P-bits.

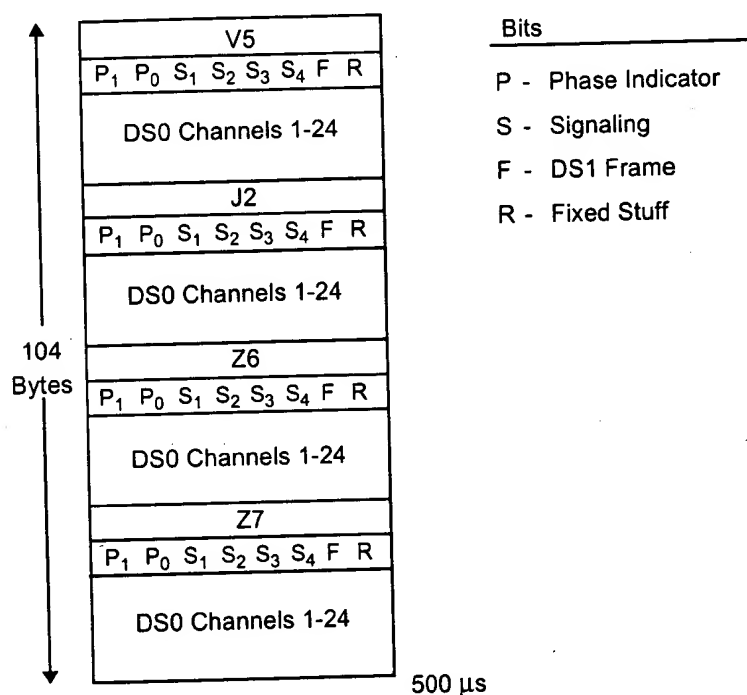


Figure 3-27. Byte-Synchronous Mapping for DS1 Payload

	Signaling												DS1 Format	
	2-State				4-State				16-State				SF	ESF
	P ₁ P ₀	S ₁	S ₂	S ₃	S ₄	S ₁	S ₂	S ₃	S ₄	S ₁	S ₂	S ₃	S ₄	F
00	A ₁	A ₂	A ₃	A ₄	A ₁	A ₂	A ₃	A ₄	A ₁	A ₂	A ₃	A ₄	F ₁	M ₁
00	A ₅	A ₆	A ₇	A ₈	A ₅	A ₆	A ₇	A ₈	A ₅	A ₆	A ₇	A ₈	S ₁	C ₁
00	A ₉	A ₁₀	A ₁₁	A ₁₂	A ₉	A ₁₀	A ₁₁	A ₁₂	A ₉	A ₁₀	A ₁₁	A ₁₂	F ₂	M ₂
00	A ₁₃	A ₁₄	A ₁₅	A ₁₆	A ₁₃	A ₁₄	A ₁₅	A ₁₆	A ₁₃	A ₁₄	A ₁₅	A ₁₆	S ₂	F ₁
00	A ₁₇	A ₁₈	A ₁₉	A ₂₀	A ₁₇	A ₁₈	A ₁₉	A ₂₀	A ₁₇	A ₁₈	A ₁₉	A ₂₀	F ₃	M ₃
00	A ₂₁	A ₂₂	A ₂₃	A ₂₄	A ₂₁	A ₂₂	A ₂₃	A ₂₄	A ₂₁	A ₂₂	A ₂₃	A ₂₄	S ₃	C ₂
01	A ₁	A ₂	A ₃	A ₄	B ₁	B ₂	B ₃	B ₄	B ₁	B ₂	B ₃	B ₄	F ₄	M ₄
01	A ₅	A ₆	A ₇	A ₈	B ₅	B ₆	B ₇	B ₈	B ₅	B ₆	B ₇	B ₈	S ₄	F ₂
01	A ₉	A ₁₀	A ₁₁	A ₁₂	B ₉	B ₁₀	B ₁₁	B ₁₂	B ₉	B ₁₀	B ₁₁	B ₁₂	F ₅	M ₅
01	A ₁₃	A ₁₄	A ₁₅	A ₁₆	B ₁₃	B ₁₄	B ₁₅	B ₁₆	B ₁₃	B ₁₄	B ₁₅	B ₁₆	S ₅	C ₃
01	A ₁₇	A ₁₈	A ₁₉	A ₂₀	B ₁₇	B ₁₈	B ₁₉	B ₂₀	B ₁₇	B ₁₈	B ₁₉	B ₂₀	F ₆	M ₆
01	A ₂₁	A ₂₂	A ₂₃	A ₂₄	B ₂₁	B ₂₂	B ₂₃	B ₂₄	B ₂₁	B ₂₂	B ₂₃	B ₂₄	S ₆	F ₃
10	A ₁	A ₂	A ₃	A ₄	A ₁	A ₂	A ₃	A ₄	C ₁	C ₂	C ₃	C ₄	F ₁	M ₇
10	A ₅	A ₆	A ₇	A ₈	A ₅	A ₆	A ₇	A ₈	C ₅	C ₆	C ₇	C ₈	S ₁	C ₄
10	A ₉	A ₁₀	A ₁₁	A ₁₂	A ₉	A ₁₀	A ₁₁	A ₁₂	C ₉	C ₁₀	C ₁₁	C ₁₂	F ₂	M ₈
10	A ₁₃	A ₁₄	A ₁₅	A ₁₆	A ₁₃	A ₁₄	A ₁₅	A ₁₆	C ₁₃	C ₁₄	C ₁₅	C ₁₆	S ₂	F ₄
10	A ₁₇	A ₁₈	A ₁₉	A ₂₀	A ₁₇	A ₁₈	A ₁₉	A ₂₀	C ₁₇	C ₁₈	C ₁₉	C ₂₀	F ₃	M ₉
10	A ₂₁	A ₂₂	A ₂₃	A ₂₄	A ₂₁	A ₂₂	A ₂₃	A ₂₄	C ₂₁	C ₂₂	C ₂₃	C ₂₄	S ₃	C ₅
11	A ₁	A ₂	A ₃	A ₄	B ₁	B ₂	B ₃	B ₄	D ₁	D ₂	D ₃	D ₄	F ₄	M ₁₀
11	A ₅	A ₆	A ₇	A ₈	B ₅	B ₆	B ₇	B ₈	D ₅	D ₆	D ₇	D ₈	S ₄	F ₅
11	A ₉	A ₁₀	A ₁₁	A ₁₂	B ₉	B ₁₀	B ₁₁	B ₁₂	D ₉	D ₁₀	D ₁₁	D ₁₂	F ₅	M ₁₁
11	A ₁₃	A ₁₄	A ₁₅	A ₁₆	B ₁₃	B ₁₄	B ₁₅	B ₁₆	D ₁₃	D ₁₄	D ₁₅	D ₁₆	S ₅	C ₆
11	A ₁₇	A ₁₈	A ₁₉	A ₂₀	B ₁₇	B ₁₈	B ₁₉	B ₂₀	D ₁₇	D ₁₈	D ₁₉	D ₂₀	F ₆	M ₁₂
11	A ₂₁	A ₂₂	A ₂₃	A ₂₄	B ₂₁	B ₂₂	B ₂₃	B ₂₄	D ₂₁	D ₂₂	D ₂₃	D ₂₄	S ₆	F ₆

Notes:

SF & ESF Format:

SF Format:

ESF Format:

F₁—F₆ = Frame Alignment Bits

S₁—S₆ = Signaling Framing Bits

C₁—C₆ = Cyclic Redundancy Check-6 Bits

M₁—M₁₂ = Data Link Bits

Figure 3-28. Byte-Synchronous DS1 Signaling and Framing Bit Assignments

The criteria in the remainder of this section apply to VT PTE with DS1 interfaces that use the byte-synchronous DS1 mapping.

R3-32 [35] VT PTE with byte-synchronous DS1 interfaces shall be capable of accepting DS1 signals using the DS1 Superframe and ESF formats defined in GR-499-CORE.

R3-33 [36] If the DS1 signal uses the ESF format and the F-bit is not used to transport the framing bits, then the following apply:

- The Cyclic Redundancy Check-6 (CRC-6) code in the received DS1 signals shall be monitored, and detected errors subsequently reported in the Data Link performance report message on the outgoing signal.
- The correct CRC-6 code shall be calculated and inserted on the outgoing DS1 signals.
- The NE shall send a performance report message every second to the sink, and receive a performance report message every second from the source.

If an NE provides DS0 rearrangement capabilities for incoming DS1 signals, then it would need to provide slip buffers so that the DS0s from different DS1s could be aligned for mapping into the VT SPE (which would be synchronized to the NE's clock, see Section 5.4). However, if the DS0 rearrangement capabilities are not provided, then the use of slip buffers could cause unnecessary DS1 frame slips in some synchronization failure scenarios. Therefore, the following objective applies, and an NE that meets the objective would use the VT Payload Pointer to frequency justify the VT SPE to the frame rate of the STS SPE (which would be synchronized to the NE's clock), as described in Section 3.5.2.

O3-34 [37] If the NE does not provide DS0 rearrangement capabilities for an incoming DS1, then it should recover clock from the DS1 and use that clock in the creation of the VT SPE.

Similarly, if the NE provides DS0 rearrangement capabilities for output DS1s, then it would need to provide slip buffers to align the DS0s from different VT SPEs, and the output DS1 would need to be synchronized to the NE's clock. However, if the DS0 rearrangement capabilities are not provided, the output DS1 could either be timed from the NE's clock, or its timing could be recovered from the payload (as is done for asynchronously mapped signals).

The signaling transport modes, DS1 framing bit transport capabilities, and pulse density assurance techniques that VT PTE must support depend on the application. Additional criteria regarding which modes and techniques must be supported appear in the criteria documents for the individual types of NEs. Shown in the following sections are the general criteria for VT PTE that supports the byte-synchronous mapping.

Section 6.2.1 contains the criteria for DS1 and VT1.5 maintenance signals (see Figures 6-9 and 6-10). In addition, SONET NEs that provide DS0 path termination or rearrangement capabilities need to support DS0 maintenance and trunk conditioning capabilities (see Section 6.2.1.6, and Figures 6-11 through 6-13).

3.4.1.1.1 DS1 Signaling and Framing Bit Transport

Two signaling transport modes are defined, as discussed below:

1. Signaling transfer mode

In the signaling transfer mode, the signaling information carried in the robbed bit positions of the DS1 signal is transferred to the S-bit positions within the VT1.5. In this mode the DS1 framing bit does not need to be transported in the F-bit position (e.g., the far-end VT PTE generates a new framing bit sequence).

CR3-35 [38] The VT PTE may be required to support the signaling transfer mode.

R3-36 [39] If the signaling transfer mode is being used, then the following apply:

- ...
 - The signaling information carried in the robbed bit positions of the DS1 signal shall be copied to the corresponding S-bit positions within the VT1.5 when the DS1 is mapped. In the VT1.5 to DS1 direction, the signaling information carried by the S-bits within the VT1.5 shall be written over the appropriate robbed bit positions of the outgoing DS1 bit stream.
- ...
 - For DS1s using the Superframe format, the robbed bit positions (from which the signaling information is copied) shall be set to '1' when the DS1 is mapped into the VT1.5.⁸
- ...
 - The phase of the S-bits shall be indicated by the P-bits in the same byte as shown in Figure 3-28 for 2-, 4-, and 16-state signaling schemes.
- ...
 - The VT PTE shall generate a new framing bit pattern for the outgoing DS1 bit stream.

2. Clear mode

In the clear mode, the VT PTE does not transfer signaling information between the S-bits and any robbed bit signaling positions (which may or may not actually contain signaling information) in the DS1 signals. It is defined for cases where (for example):

8. For DS1s using the ESF format, the VT PTE may either set the robbed bit positions to 1, or leave them set to their existing values.

- The DS1 framing bit is carried in the F-bit to identify the phase of the robbed bit signaling information contained in the mapped DS1.
- The signaling is carried in a separate channel (e.g., the Common Signaling Channel in some IDLC systems), so robbed bit signaling is not used in the DS1.

CR3-37 [40] The VT PTE may be required to support the clear mode.

CR3-38 [41] If the clear mode is being used, then the VT PTE may be required to be user-provisionable (on a per-DS1 basis) to transport the DS1 frame bit in the F-bit position of the VT1.5.

R3-39 [42] If the clear mode is being used and the DS1 frame bit is carried in the F-bit, then the framing bits in the incoming DS1 shall be placed in the transmitted F-bits, and the received F-bits shall be placed in the outgoing DS1 signal (with no change in phase relative to the DS0 channels). The phase of the F-bit shall be indicated by the P-bits in the same byte as shown in Figure 3-28 for the Superframe and ESF formats.

R3-40 [43] If the clear mode is being used and the DS1 frame bit is not carried in the F-bit (i.e., the F-bit is undefined), then the VT PTE shall generate a new framing bit pattern for the outgoing DS1 bit stream.

In addition to the two signaling transport modes defined above for use on a per-DS1 basis, in some applications it may be desirable for the VT PTE to perform, or not perform, signaling transfer on a per-DS0 basis. In some applications, this capability may only need to be provided on a static (i.e., user-provisionable) basis.

CR3-41 [44] The VT PTE may be required to be user-provisionable on a per-DS0 basis to either perform or not perform signaling transfer (i.e., to provide signaling transfer/clear DS0 transport).

In other applications, signaling transfer/clear DS0 transport may need to be provided on a dynamic basis. To accomplish this, a scheme has been defined in which the S-bits are set to a specific value for the DS0s for which signaling transfer is not to be performed, and all other values in the S-bits indicate that signaling transfer is to be performed. In this scheme, the assignment of a particular DS0 as a clear DS0 is made by two NEs (e.g., an IDT and an RDT in an IDLC system) that are able to communicate that assignment with each other and set the S-bits accordingly. Intermediate NEs then need to be able to adjust their operation based on the code in the S-bits. This scheme avoids the difficulty of provisioning each DS0 channel at a DS1 interface and also allows dynamic signaling transfer/clear DS0 assignment.

- CR3-42** [45] The VT PTE may be required to be user-provisionable (on either a per-DS0 or per-DS1 basis) to perform dynamic signaling transfer/clear DS0 transport.
- R3-43** [46] If dynamic signaling transfer/clear DS0 transport is being used, then the code ABCD=1001 on the S-bits shall be interpreted to mean that signaling transfer is not to be performed on that DS0 channel in either direction of transmission. Signaling transfer shall be performed for DS0 channels whose associated S-bits do not contain the code ABCD=1001.

Note that the ABCD code used to indicate that signaling transfer is not to be performed on a particular DS0 (in both directions) only appears in the S-bits in the SONET signal. It does not appear in any form in the DS1 signal. Therefore, this scheme will only support a single byte-synchronous SONET to DS1 and DS1 to byte-synchronous SONET conversion between the two NEs that make the signaling transfer/clear DS0 assignments and set the S-bits to the appropriate codes.

3.4.1.1.2 *Line Codes and Pulse Density Assurance*

The following criteria on DS1 line codes and pulse density assurance techniques are applicable:

- R3-44** [47] The VT PTE shall accommodate both the Alternate Mark Inversion (AMI) line code, and the Bipolar with Eight Zero Substitution (B8ZS) line code.

Although a DS1 signal that uses the AMI line code and appears at a DS1 interface is required to have no more than 15 consecutive zeros and at least N ones in each and every time window of 8 (N+1) digital time slots (see GR-499-CORE), the AMI line code itself does not provide any form of pulse density assurance. Several methods of assuring the necessary pulse density are available, including the B8ZS line code, Zero Byte Time Slot Interchange (ZBTISI), and Zero Code Suppression (ZCS).⁹ In most applications it is expected that SONET NEs will use B8ZS unless they are connected to NEs that do not support that line code. If a SONET NE is simply providing transport for DS1s created by other NEs, then it is the responsibility of those source NEs to perform ZBTISI or ZCS (or to support B8ZS). However, if a SONET NE provides DS0 path termination or DS0 rearrangement capabilities, then it is a DS1 source (even though the DS1 that it creates may not appear at a DS1 interface until the VT1.5 is terminated at another SONET NE). Therefore, a source SONET NE may need to support ZBTISI or ZCS if the DS1s that it creates are expected to be received by NEs that do not support B8ZS.

9. Criteria on ZBTISI are contained in GR-499-CORE. ZCS, which is also called "bit stuffing" in some documents, involves the insertion (at the DS1 source) of a '1' in bit 7 of any all-zeros DS0 byte.

CR3-45 [48] VT PTE that supports DS0 path terminations or DS0 rearrangement capabilities may be required to support ZBTISI for pulse density assurance.

CR3-46 [49] VT PTE that supports DS0 path terminations or DS0 rearrangement capabilities may be required to support ZCS for pulse density assurance.

In addition, in some applications it may be necessary for VT PTE to support ZBTISI even if it does not support DS0 path terminations or DS0 rearrangement capabilities. For example, a SONET NE that supports ZBTISI could be used to convert DS1s that are transported (via SONET signals) between a DS1 network that uses ZBTISI and one that uses B8ZS.

CR3-47 [50] VT PTE may be required to support ZBTISI.

R3-48 [51] If ZBTISI is supported, then the ZBTISI algorithm and the ESF data link as described in GR-499-CORE shall be used.

R3-49 [52] The choice between AMI or B8ZS, and of ZBTISI or ZCS (if provided), shall be provisionable by the user on a per-DS1 interface basis.

3.4.1.2 Asynchronous Mapping for DS1

An asynchronous mapping of a DS1 into the payload capacity of a VT1.5 SPE is defined for clear-channel transport of DS1 signals that meet the DSX-1 requirements in GR-499-CORE. If the asynchronous DS1 mapping is supported, then the following criteria are applicable.

R3-50 [53] The asynchronous mapping of a DS1 into a VT1.5 SPE shall be as shown in Figure 3-29.

The asynchronous DS1 mapping contains 771 information (I) bits, 6 stuff control (C) bits, 2 stuff opportunity (S) bits, and 8 overhead communication channel (O) bits in each VT1.5 SPE. The remaining 13 bits of the VT1.5 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future communication purposes. The values contained in the R- and O-bits are currently undefined.

R3-51 [54] In each VT1.5 SPE, two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.

The value contained in the S-bits when they are stuff bits is undefined.

- R3-52** [55] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single bit errors in the C-bits.
- R3-53** [56] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS1 jitter transfer mask shown in Figure 5-26, the output jitter is less than 0.7 Unit Intervals peak-to-peak (UI_{pp}), assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
- R3-54** [903] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS1 jitter transfer mask shown in Figure 5-26, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than that same DS1 jitter transfer mask.
- R3-55** [57] The DS1 interface shall accommodate both the AMI line code (assuming the DS1 source meets the zeros constraints in GR-499-CORE, see Section 3.4.1.1.2) and the B8ZS line code.
- R3-56** [58] The choice of AMI or B8ZS shall be provisionable by the user on a per-DS1 interface basis.

Section 6.2.1 contains the criteria for DS1 and VT1.5 maintenance signals (see Figure 6-7).

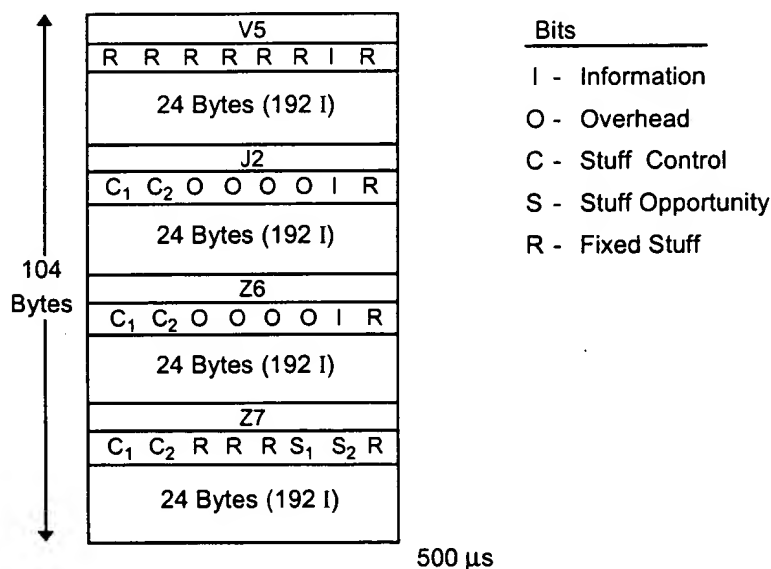


Figure 3-29. Asynchronous Mapping for DS1 Payload

3.4.1.3 Asynchronous Mapping for DS1C

An asynchronous mapping of a DS1C into the payload capacity of a VT3 SPE is defined for clear-channel transport of DS1C signals that meet the DSX-1C requirements in GR-499-CORE. If the asynchronous DS1C mapping is supported, then the following criteria are applicable.

- R3-57** [59] The asynchronous mapping of a DS1C into a VT3 SPE shall be as shown in Figure 3-30.

The asynchronous DS1C mapping contains 1574 information (I) bits, 12 stuff control (C) bits, 4 stuff opportunity (S) bits, and 16 overhead communication channel (O) bits in each VT3 SPE. The remaining 58 bits of the VT3 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future communication purposes. The values contained in the R- and O-bits are currently undefined.

- R3-58** [60] Twice in each VT3 SPE, the two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.

The value contained in the S-bits when they are stuff bits is undefined.

- R3-59** [61] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single bit errors in the C-bits.
- R3-60** [62] The stuffing mechanism that generates the C-bits shall be chosen so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 350 Hz, the output jitter is less than $1.0 UI_{pp}$ and $0.3 UI_{rms}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
- R3-61** [904] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 350 Hz, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than the DS1C jitter transfer mask in Section 7.3.2 of GR-499-CORE.
- R3-62** [63] The DS1C interface shall accommodate both the AMI line code (assuming the DS1C source meets the ones density criteria from GR-499-CORE of at least 12.5% ones over any 150 consecutive bits), and the B8ZS line code.

R3-63 [64] The choice of AMI or B8ZS shall be provisionable by the user on a per-DS1C interface basis.

Section 6.2.1 contains the criteria for DS1C and VT3 maintenance signals (see Figure 6-7).

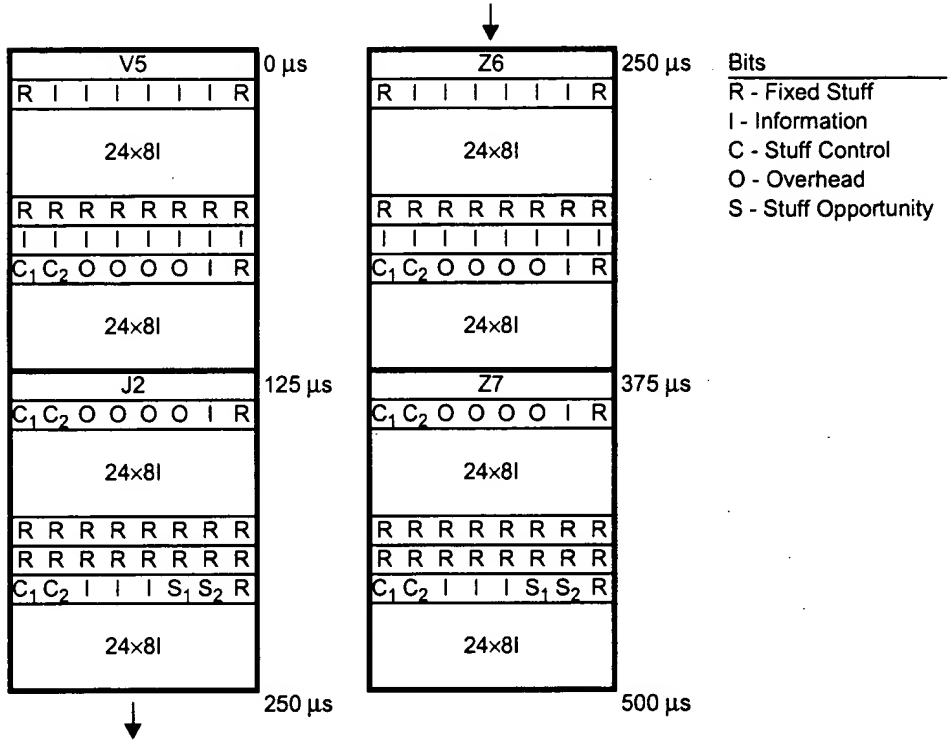


Figure 3-30. Asynchronous Mapping for DS1C Payload

3.4.1.4 Asynchronous Mapping for DS2

An asynchronous mapping of a DS2 into the payload capacity of a VT6 SPE is defined for clear-channel transport of DS2 signals that meet the DSX-2 requirements in GR-499-CORE. If the asynchronous DS2 mapping is supported, then the following criteria are applicable.

R3-64 [65] The asynchronous mapping of a DS2 into a VT6 SPE shall be as shown in Figure 3-31.

The asynchronous DS2 mapping contains 3152 information (I) bits, 24 stuff control (C) bits, 8 stuff opportunity (S) bits, and 32 overhead communication channel (O) bits in each

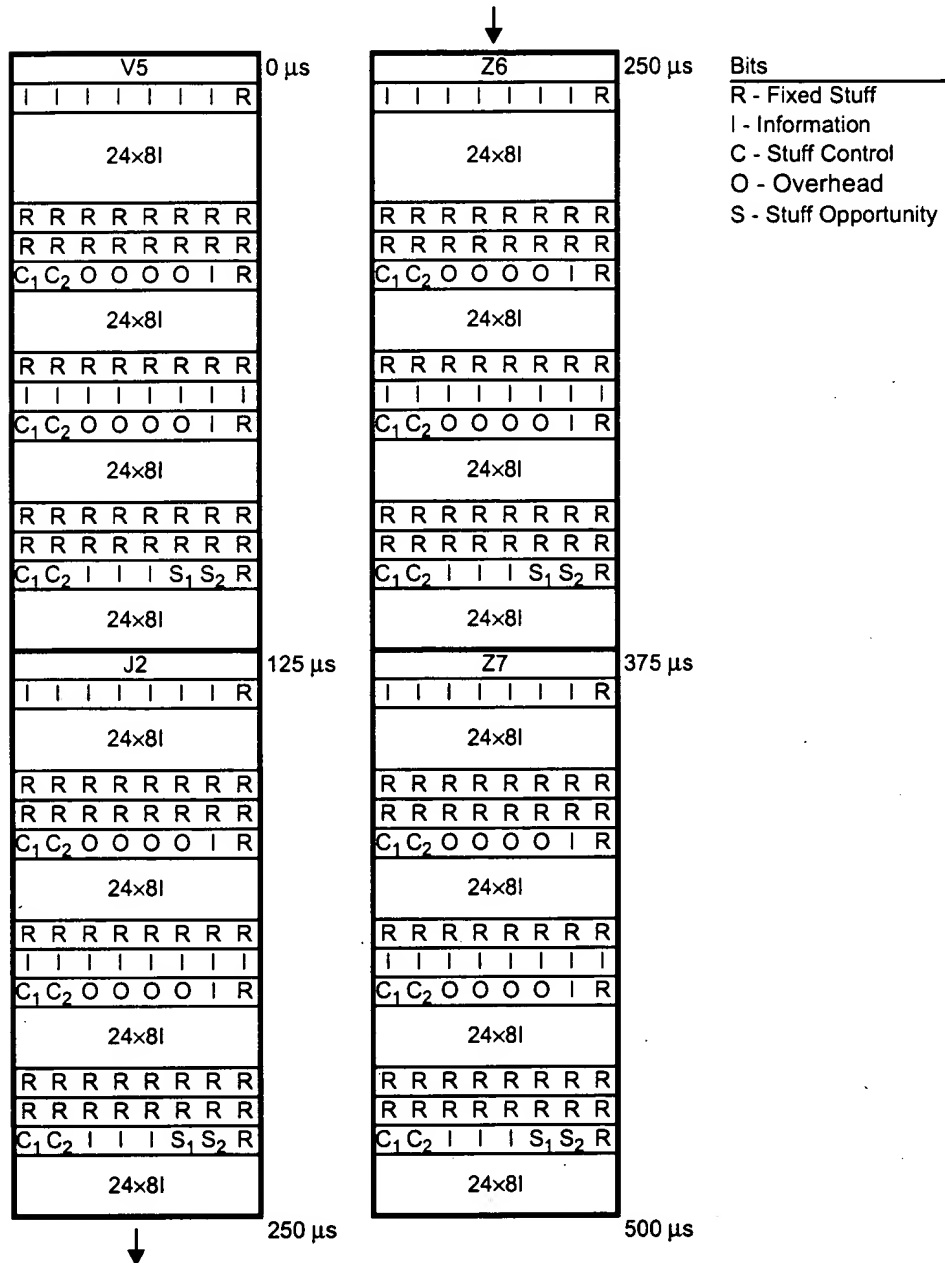
VT6 SPE. The remaining 176 bits of the VT6 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future overhead communication purposes. The values contained in the R- and O-bits are currently undefined (see Section 3.2).

- R3-65** [66] Four times in each VT6 SPE, the two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.

The value contained in the S-bits when they are stuff bits is undefined.

- R3-66** [67] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single-bit errors in the C-bits.
- R3-67** [68] The stuffing mechanism that generates the C-bits shall be chosen so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 500 Hz, the output jitter is less than $1.0 UI_{pp}$ and $0.3 UI_{rms}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
- R3-68** [905] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 500 Hz, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than the DS2 jitter transfer mask in Section 7.3.2 of GR-499-CORE.

Section 6.2.1 contains the criteria for DS2 and VT6 maintenance signals (see Figure 6-7).



3.4.2 STS-1 Mappings

This section describes mappings that occupy the entire payload capacity of an STS-1 SPE (see Figure 3-4).

3.4.2.1 Asynchronous Mapping for DS3

An asynchronous mapping for a DS3 into the payload capacity of an STS-1 SPE is defined for clear-channel transport of DS3 signals that meet the DSX-3 requirements in GR-499-CORE. If the asynchronous DS3 mapping is supported, then the following criteria are applicable.

- R3-69** [69] The asynchronous mapping for a DS3 into an STS-1 SPE shall be as shown in Figure 3-32.

The asynchronous DS3 mapping consists of 9 subframes every 125 μ s. Each subframe contains 621 information (I) bits, a set of 5 stuff control (C) bits, 1 stuff opportunity (S) bit, and 2 overhead communication channel (O) bits. The remaining bits of the STS-1 Payload Capacity are fixed stuff (R) bits. The O-bits are reserved for future overhead communication purposes. The values contained in the R- and O-bits are currently undefined.

- R3-70** [70] In each subframe, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.

The value contained in the S-bit when it is a stuff bit is undefined.

- R3-71** [71] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.
- R3-72** [906] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS3 jitter transfer mask shown in Figure 5-26, the output jitter is less than 0.4 UI_{pp}, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
- R3-73** [907] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS3 jitter transfer mask shown in Figure 5-26, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than that same DS3 jitter transfer mask.

Section 6.2.1 contains the criteria for DS3 and STS-1 path maintenance signals (see Figure 6-5).

← 28 Bytes →					← 28 Bytes →					← 28 Bytes →				
STS POH	R	R	C1	25 I	Fixed Stuff	R	C2	I	25 I	Fixed Stuff	R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I

Bytes

I =

i	i	i	i	i	i	i	i
---	---	---	---	---	---	---	---

R =

r	r	r	r	r	r	r	r
---	---	---	---	---	---	---	---

C1 =

r	r	c	i	i	i	i	i
---	---	---	---	---	---	---	---

C2 =

c	c	r	r	r	r	r	r
---	---	---	---	---	---	---	---

C3 =

c	c	r	r	o	o	r	s
---	---	---	---	---	---	---	---

bits

i: information (payload) bit

r: fixed stuff bit

c: stuff control bit

s: stuff opportunity bit

o: overhead communications channel bit

Figure 3-32. Asynchronous Mapping for DS3 Payload

3.4.2.2 Asynchronous Transfer Mode (ATM) Mapping for B-ISDN Applications

A method for mapping ATM cells (each of which consists of a 5-byte cell header and a 48-byte payload) into the payload capacity of an STS-1 SPE is defined. If this mapping is supported, then the following requirement is applicable.

- R3-74** [72] ATM cells shall be mapped into the STS-1 Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-1 SPE. The entire STS-1 Payload Capacity (i.e., 84 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 48.384 Mb/s.

Because the STS-1 Payload Capacity is not an integer multiple of the 53-byte ATM cell length, some cells will cross an SPE boundary.

Cell payload scrambling is used to provide security against payload information replicating the frame synchronous scrambling sequence (or its inverse) used at the SONET Section layer. Details of the cell payload scrambler are contained in TR-NWT-001112, *Broadband-ISDN User to Network Interface and Network Node Interface Physical Layer Generic Criteria*.

3.4.2.3 HDLC-Over-SONET Mapping

This section contains criteria related to a mapping for HDLC-framed signals into the payload capacity of an STS-1 SPE. Equivalent mappings can be used for the transport of HDLC-framed signals in other sizes of STS SPEs such as STS-3c or STS-12c SPEs. In such cases, the text and criteria shown below apply, with the only changes being the descriptions of the number of columns in the payload capacity and the maximum transfer capacity for the HDLC-framed signals.

Scrambling of HDLC-framed signals is necessary to provide security against payload information replicating (for example) the frame synchronous scrambling sequence used at the SONET Section layer. As described below, a self-synchronous scrambler with a generator polynomial of $x^{43}+1$ is used. Note that no initial seed is specified for this scrambler, and therefore the first 43 bits that are transmitted following a startup or reframe operation cannot be descrambled correctly.

R3-75 [1044] The following apply if the HDLC-over-SONET mapping is supported.

- ... • The HDLC-framed signal shall be mapped into the STS-1 Payload Capacity by aligning the byte structure of every frame with the byte structure of the STS-1 SPE.
- ... • HDLC flags (i.e., '01111110' bytes) shall be used for interframe fill to account for the variable nature of the arrival of the HDLC frames.
- ... • The entire STS-1 Payload Capacity (i.e., 84 columns) shall be filled with HDLC frames and HDLC flags (as necessary).
- ... • The HDLC-framed signal plus the interframe fill shall be scrambled before it is inserted into the STS-1 Payload Capacity. In the reverse operation, after the STS-1 path is terminated the payload shall be descrambled before it is passed to the HDLC layer.
- ... • A self-synchronizing scrambler with a generator polynomial of $x^{43}+1$ shall be used.

- ...
 - The most significant bit of each byte of the HDLC-framed signal and interframe fill (i.e., the bit that will be placed into bit 1 of a byte in the STS-1 Payload Capacity, see Figure 3-2) shall enter the scrambler first, followed by the next most significant bit of that byte, etc.
- ...
 - The scrambler shall run continuously (e.g., it shall not be reset for each SONET or HDLC frame).

Filling the entire STS-1 Payload Capacity yields a maximum transfer capacity of 48.384 Mb/s for HDLC-framed signals carried in a nominal rate STS-1 SPE. Because HDLC frames are of variable length, a frame may cross an SPE boundary.

Typically, HDLC frames include a CRC for error detection purposes. Two CRCs that are often used are a CRC-16 and a CRC-32. The use of the CRC-32 is preferred when the HDLC frames are carried over SONET, as that reduces the impact of the error multiplication effect caused by the $x^{43}+1$ scrambler.

- O3-76** **[1045]** If the HDLC-over-SONET mapping is supported and a Cyclic Redundancy Check is applied over the HDLC payload signal, a CRC-32 should be used.

3.4.3 Super Rate Payloads

This section describes the mappings of Super Rate payloads into the payload capacities of STS-Nc SPEs.

3.4.3.1 Mappings into an STS-3c SPE

The STS-3c SPE is a 261-column by 9-row structure. The first column contains the STS POH bytes, and the remaining 260 columns are the STS-3c Payload Capacity.

3.4.3.1.1 Asynchronous Mapping for DS4NA

An asynchronous mapping of a DS4NA (which has a nominal bit rate of 139.264 Mb/s) into the STS-3c Payload Capacity is defined for clear-channel transport of DS4NA signals that meet the DSX-4NA requirements in GR-499-CORE. If the asynchronous DS4NA mapping is supported, then the following criteria are applicable.

- R3-77** **[73]** The asynchronous mapping of a DS4NA into an STS-3c SPE shall be as shown in Figure 3-33.

Each 260-byte row of the STS-3c Payload Capacity is divided into 20 blocks of 13 bytes each (see Figure 3-33 for the detailed structure of the blocks), and contains 1934

information (I) bits, a set of 5 stuff control (C) bits, 1 stuff opportunity (S) bit, 10 overhead communication channel (O) bits, and 130 fixed stuff (R) bits. The O-bits are reserved for future overhead communications purposes. The values contained in the R- and O-bits are currently undefined.

- R3-78** [74] In each row, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.

The value contained in the S-bit when it is a stuff bit is undefined.

- R3-79** [75] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.

Section 6.2.1 contains the criteria for DS4NA and STS-3c path maintenance signals.

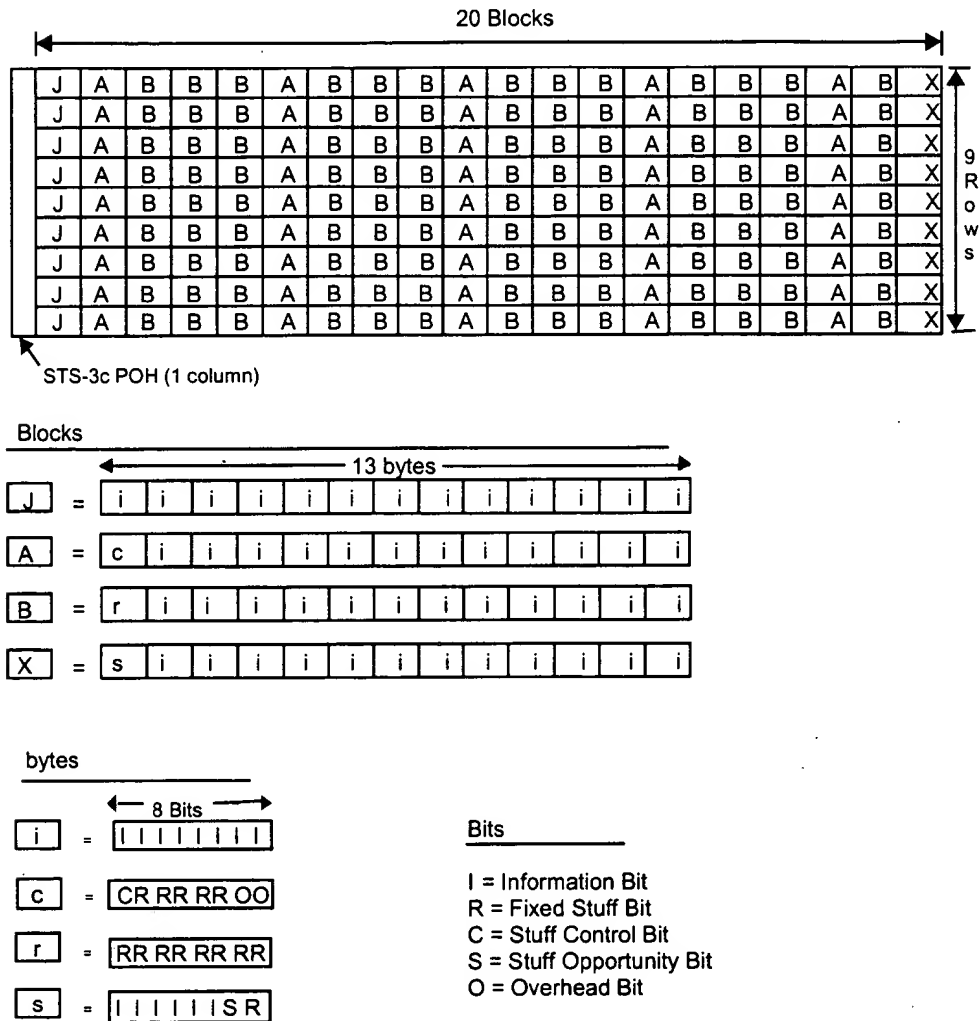


Figure 3-33. Asynchronous Mapping for DS4NA Payload

3.4.3.1.2 *Asynchronous Mapping for Fiber Distributed Data Interface (FDDI) at 125 Mb/s*

An asynchronous mapping of an FDDI signal into the STS-3c Payload Capacity is defined for clear-channel transport of FDDI signals. If this mapping is supported, then the following criteria are applicable.

- R3-80** [76] The asynchronous mapping for a 125-Mb/s FDDI signal into an STS-3c SPE shall be as shown in Figure 3-34.

Each 260-byte row of the STS-3c Payload Capacity is divided into 20 blocks of 13 bytes each (see Figure 3-34 for the detailed structure of the blocks), and contains 1736 or 1735 information (I) bits, one set of 5 stuff control (C) bits, 1 stuff opportunity (S) bit, 2 overhead communication channel (O) bits, and 336 or 337 fixed stuff (R) bits. The O-bits are reserved for future overhead communications purposes. The values contained in the R- and O-bits are currently undefined.

- R3-81** [77] In each row, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.

The value contained in the S-bit when it is a stuff bit is undefined.

- R3-82** [78] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.

Section 6.2.1 contains the criteria for FDDI and STS-3c path maintenance signals.

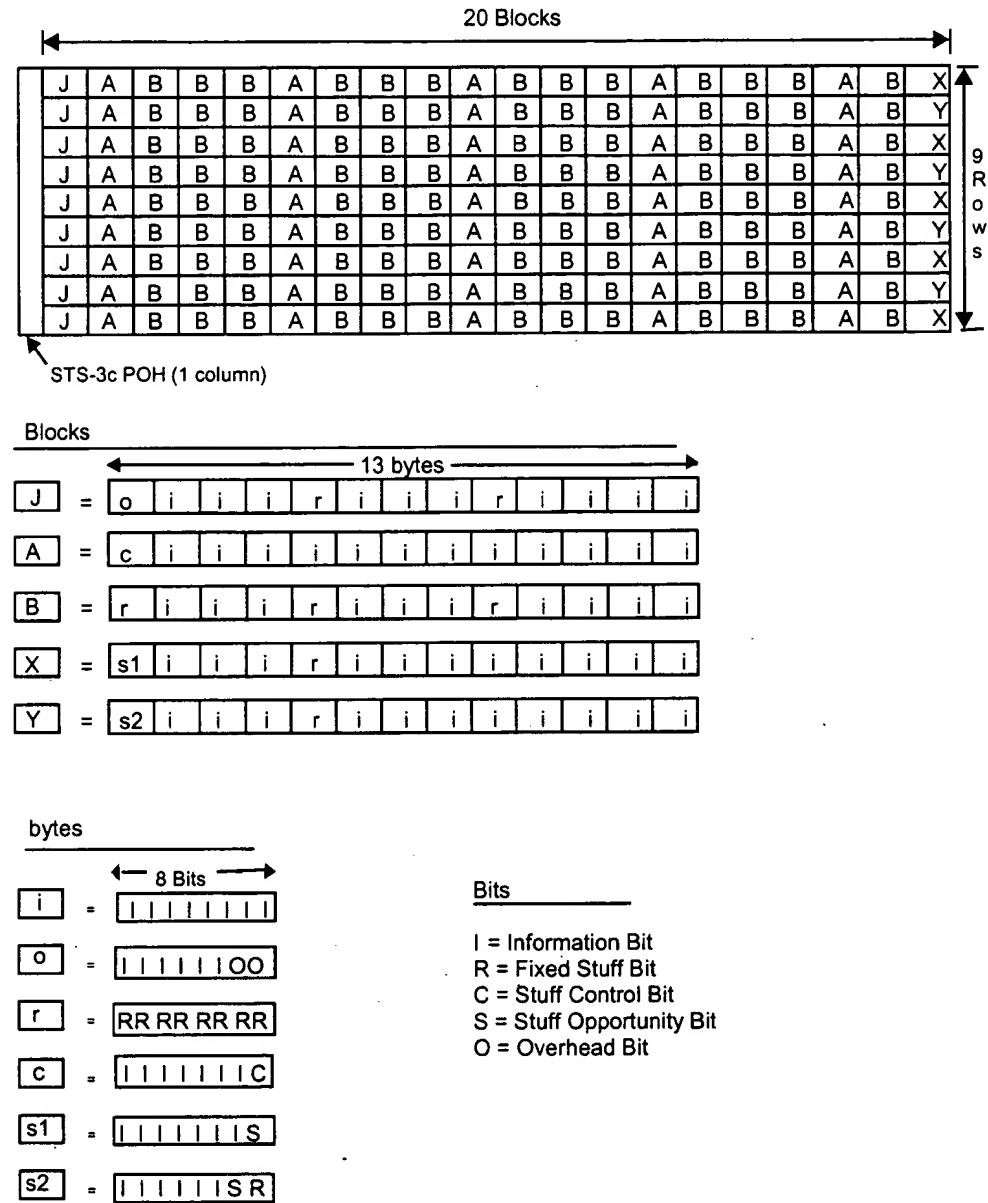


Figure 3-34. Asynchronous Mapping for FDDI

3.4.3.1.3 *Asynchronous Transfer Mode (ATM) Mapping for B-ISDN Applications*

A method for mapping ATM cells (each of which consists of a 5-byte cell header and a 48-byte payload) into the payload capacity of an STS-3c SPE is defined. If this mapping is supported, then the following requirement is applicable.

- R3-83** [79] ATM cells shall be mapped into the STS-3c Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-3c SPE. The entire STS-3c Payload Capacity (i.e., 260 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 149.760 Mb/s.

Because the STS-3c Payload Capacity is not an integer multiple of the 53-byte ATM cell length, some cells will cross an SPE boundary.

Cell payload scrambling is used to provide security against payload information replicating the frame synchronous scrambling sequence (or its inverse) used at the SONET Section layer. Details of the cell payload scrambler are contained in TR-NWT-001112.

3.4.3.1.4 *DQDB Metropolitan Area Network (MAN) Mapping*

A mapping for DQDB into an STS-3c SPE is defined in this section. In this mapping, DQDB layer slots are transported in the payload capacity of an STS-3c SPE. In addition, the mapping uses three of the STS POH bytes, as follows:

- The F2 (User Channel) and Z3 (Growth) bytes are used to carry the DQDB Layer Management (M1 and M2) bytes.
- Bits 1 and 2 of the H4 (Indicator) byte are used to carry the Link Status Signal (LSS).
- Bits 3 through 8 of the H4 byte are used to indicate the offset from the H4 byte to the beginning of the next 53-byte DQDB slot.

If the DQDB mapping is supported, then the following criteria apply.

- R3-84** [80] DQDB slots shall be mapped into the STS-3c Payload Capacity by aligning the byte structure of every slot with the byte structure of the STS-3c SPE. The entire STS-3c Payload Capacity (i.e., 260 columns) shall be filled with slots, yielding a transfer capacity for DQDB slots of 149.760 Mb/s.

- R3-85** [81] Bits 3 through 8 of the H4 byte shall contain a binary number in the range from '000000' (0) to '110100' (52) that indicates the offset between the H4 byte and the boundary of the first DQDB slot following the H4 byte.

The use of H4 to indicate the offset to the next DQDB slot is illustrated in Figure 3-35, and the bit assignments for the H4 byte are shown in Figure 3-36.

Because the STS-3c Payload Capacity is not an integer multiple of the DQDB slot length, some slots will cross the SPE boundary. Also, cell payload scrambling is used to provide security against payload information replicating the frame synchronous scrambling sequence (or its inverse) used at the SONET Section layer.

Information about the status of the transmission link between the two adjacent Physical Layer Convergence Procedure (PLCP) entities connected to the ends of the transmission link is communicated by the LSS.

R3-86 [82] Bits 1 and 2 of the H4 byte shall be used to carry the LSS.

R3-87 [83] The F2 and Z3 bytes of the STS POH shall carry the DQDB M1 and M2 bytes, respectively.

The DQDB M1 and M2 bytes transport the DQDB layer node management information. These management bytes are generated at the head of a bus and are employed by the DQDB Layer Management Protocol Entity as IEEE P802.6/D14, *Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN)*, describes.

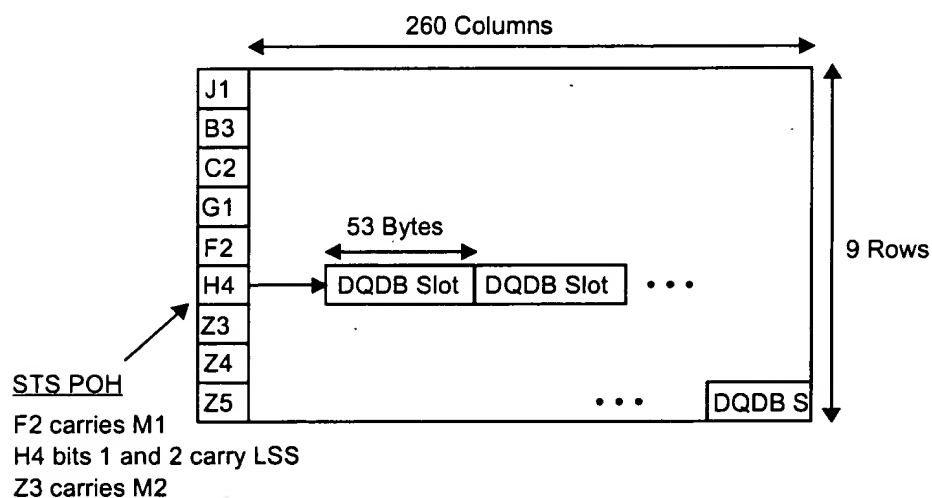


Figure 3-35. DQDB Mapping into an STS-3c SPE

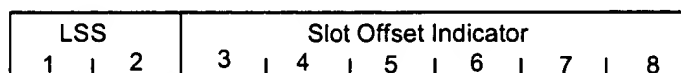


Figure 3-36. Bit Allocation for H4 Byte in DQDB Mapping

3.4.3.1.5 *HDLC-Over-SONET Mapping*

As noted in Section 3.4.2.3, the HDLC-over-SONET mapping described in that section can also be used for the transport of HDLC-framed signals in STS-3c SPEs. If such a mapping is supported, the criteria in Section 3.4.2.3 apply, with "STS-1" replaced by "STS-3c".

As was the case for the STS-1 HDLC-over-SONET mapping, the entire STS-3c Payload Capacity (i.e., 260 columns) is filled with HDLC frames and HDLC flags (as necessary). This results in a maximum transfer capacity of 149.760 Mb/s for HDLC-framed signals carried in a nominal rate STS-3c SPE.

3.4.3.2 *Mappings into STS-12c SPEs*

The STS-12c SPE is a 1044-column by 9-row structure. The first column contains the STS POH bytes. In all of the currently defined mapping, the STS POH is followed by three columns of fixed stuff bytes, and the remaining 1040 columns are the STS-12c Payload Capacity.

3.4.3.2.1 *Asynchronous Transfer Mode (ATM) Mapping for B-ISDN Applications*

A method for mapping ATM cells (each of which consists of a 5-byte cell header and a 48-byte payload) into the payload capacity of an STS-12c SPE is defined. If this mapping is supported, then the following requirement is applicable.

- R3-88** **[84]** ATM cells shall be mapped into the STS-12c Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-12c SPE. The entire STS-12c Payload Capacity (i.e., 1040 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 599.040 Mb/s.

Because the STS-12c Payload Capacity is not an integer multiple of the 53-byte ATM cell length, some cells will cross an SPE boundary.

Cell payload scrambling is used to provide security against payload information replicating the frame synchronous scrambling sequence (or its inverse) used at the SONET Section layer. Details of the cell payload scrambler are contained in TR-NWT-001112.

3.4.3.2.2 *HDLC-Over-SONET Mapping*

As noted in Section 3.4.2.3, the HDLC-over-SONET mapping described in that section can also be used for the transport of HDLC-framed signals in STS-12c SPEs. If such a mapping is supported, the criteria in Section 3.4.2.3 apply, with "STS-1" replaced by "STS-12c".

As was the case for the STS-1 HDLC-over-SONET mapping, the entire STS-12c Payload Capacity (i.e., 1040 columns) is filled with HDLC frames and HDLC flags (as necessary). This results in a maximum transfer capacity of 599.040 Mb/s for HDLC-framed signals carried in a nominal rate STS-12c SPE.

3.5 Payload Pointers

3.5.1 STS Payload Pointer

The STS Payload Pointer provides a method of allowing flexible and dynamic alignment of the STS SPE within the STS Envelope Capacity, independent of the actual contents of the SPE. Dynamic alignment means that the STS SPE is allowed to float within the STS Envelope Capacity. Thus, the pointer is able to accommodate differences not only in the phases of the STS SPE and the Transport Overhead, but in the frame rates as well.

The Payload Pointer is contained in the H1 and H2 bytes of the Line Overhead, and designates the location of the byte where the STS SPE begins. These two bytes can be viewed as one word, as shown in Figure 3-37. Bits 1 through 4 of the pointer word carry the New Data Flag, and bits 7 through 16 carry the pointer value. Bits 5 and 6 of the STS pointer word are undefined, and therefore the criteria in Section 3.2 are applicable.

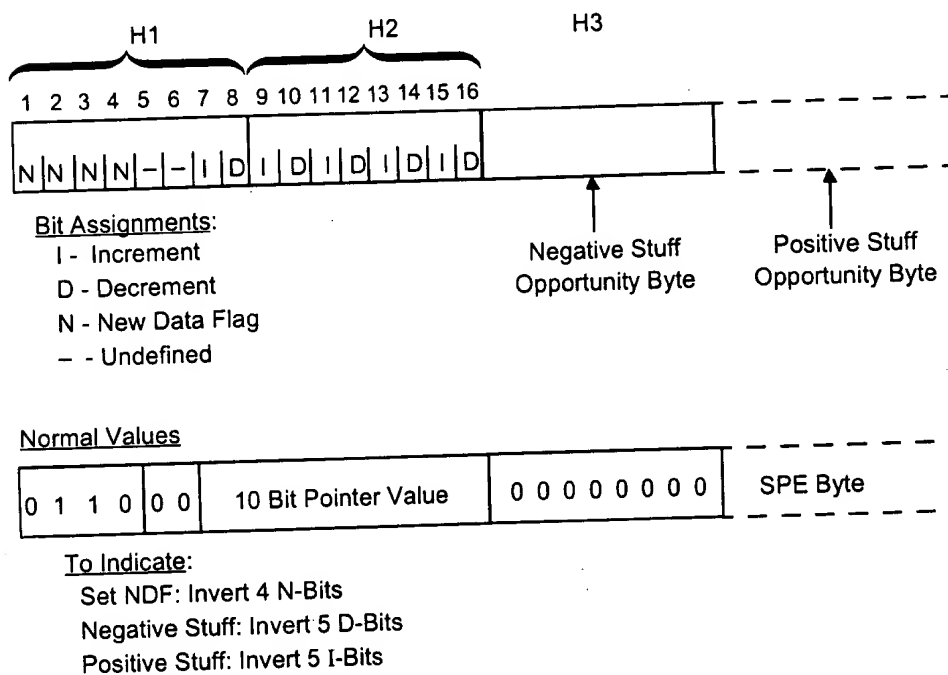


Figure 3-37. STS Payload Pointer (H1, H2) Coding

3.5.1.1 Pointer Value

Bits 7 through 16 of the pointer word carry the pointer value, which indicates the offset between the pointer word and the first byte of the STS SPE (i.e., the J1 byte). The Transport Overhead bytes are not counted in the offset. For example, a pointer value of 0 indicates that the STS SPE starts in the byte location that immediately follows the H3 byte, while an offset of 87 indicates that it starts immediately after the K2 byte location.

R3-89 [85] The pointer value shall be a binary number with a range of 0 to 782, and shall indicate the offset between the pointer word and the first byte of the STS SPE (as shown in Figure 3-38).

Note that in the case of an STS-Nc SPE, there are at least two ways to view the relationship between the pointer value and the offset from the pointer word to the J1 byte. For example, it could be considered that there is a one-to-one correspondence, and that only the STS-Nc envelope capacity bytes that are associated with the first STS-1 of the STS-Nc are counted in determining the offset. Alternatively, all of the bytes in the STS-Nc envelope capacity could be counted in determining the offset, and the NE could then transmit a pointer value

The diagram illustrates the structure of an STS-1 Frame, which is 250 μs in duration. The frame is divided into three main sections: Transport Overhead (0 to 3 μs), Negative Stuff Opportunity (3 to 86 μs), and Positive Stuff Opportunity (86 to 250 μs).

Transport Overhead (0 to 3 μs): This section contains the frame header and control fields. It is divided into three columns labeled 1, 2, and 3. The first column (1) contains H1, the second column (2) contains H2, and the third column (3) contains H3. The first row of this section is labeled 0 μs. The second row is labeled 86 μs. The third row is labeled 87 μs. The fourth row is labeled 88 μs. The fifth row is labeled 522 μs. The sixth row is labeled 521 μs. The seventh row is labeled 781 μs. The eighth row is labeled 782 μs. The ninth row is labeled 0 μs. The tenth row is labeled 86 μs. The eleventh row is labeled 87 μs. The twelfth row is labeled 88 μs. The thirteenth row is labeled 522 μs. The fourteenth row is labeled 521 μs. The fifteenth row is labeled 781 μs. The sixteenth row is labeled 782 μs.

Negative Stuff Opportunity (3 to 86 μs): This section is indicated by a dashed line and contains the payload data. It is divided into three columns labeled 1, 2, and 3. The first column (1) contains H1, the second column (2) contains H2, and the third column (3) contains H3. The first row of this section is labeled 0 μs. The second row is labeled 86 μs. The third row is labeled 87 μs. The fourth row is labeled 88 μs. The fifth row is labeled 522 μs. The sixth row is labeled 521 μs. The seventh row is labeled 781 μs. The eighth row is labeled 782 μs. The ninth row is labeled 0 μs. The tenth row is labeled 86 μs. The eleventh row is labeled 87 μs. The twelfth row is labeled 88 μs. The thirteenth row is labeled 522 μs. The fourteenth row is labeled 521 μs. The fifteenth row is labeled 781 μs. The sixteenth row is labeled 782 μs.

Positive Stuff Opportunity (86 to 250 μs): This section is indicated by a dashed line and contains the payload data. It is divided into three columns labeled 1, 2, and 3. The first column (1) contains H1, the second column (2) contains H2, and the third column (3) contains H3. The first row of this section is labeled 0 μs. The second row is labeled 86 μs. The third row is labeled 87 μs. The fourth row is labeled 88 μs. The fifth row is labeled 522 μs. The sixth row is labeled 521 μs. The seventh row is labeled 781 μs. The eighth row is labeled 782 μs. The ninth row is labeled 0 μs. The tenth row is labeled 86 μs. The eleventh row is labeled 87 μs. The twelfth row is labeled 88 μs. The thirteenth row is labeled 522 μs. The fourteenth row is labeled 521 μs. The fifteenth row is labeled 781 μs. The sixteenth row is labeled 782 μs.

3.5.1.2 STS Frequency Justification

When the frame rate of the STS SPE is less than that of the Transport Overhead, the alignment of the SPE is periodically slipped back in time (using a positive stuff byte) and the pointer value is incremented by one. Similarly, when the frame rate of the STS SPE is greater than that of the Transport Overhead, the alignment of the SPE is periodically advanced in time (using a “negative stuff” byte) and the pointer value is decremented by one. In both cases, subsequent pointers contain the new offset.

- R3-90** [86] When there is a frequency offset between the frame rate of the Transport Overhead and that of the STS SPE, the pointer value shall be incremented or decremented as needed (although see rule 7 of **R3-100** [96v2]), accompanied by a corresponding positive or negative stuff byte.
- R3-91** [87] A pointer increment operation shall be indicated by inverting bits 7, 9, 11, 13, and 15 (the I-bits) of the pointer word. The positive stuff byte shall appear immediately after the H3 byte in the frame containing the inverted I-bits, as shown in Figure 3-39.

The value contained in the positive stuff byte is undefined.

- R3-92** [88] A pointer decrement operation shall be indicated by inverting bits 8, 10, 12, 14, and 16 (the D-bits) of the pointer word. The H3 byte shall be used as the negative stuff byte, (i.e., it is used to carry an SPE byte in the frame containing the inverted D-bits), as shown in Figure 3-40.

The following criteria are applicable to the receiving STS pointer processor, and provide protection against errors on the I- and D-bits. The method in the objective is an extension of the majority vote method for I- and D-bits separately, and enhances performance during error bursts.

- R3-93** [89] If **O3-94** [90] (the "8 of 10" objective) is not met, then the increment decision shall be made by a majority vote of the I-bits, and the decrement decision shall be made by a majority vote of the D-bits.¹⁰
- O3-94** [90] The increment/decrement decision should be made at the receiver by a match of 8 or more of the 10 I- and D-bits to either the increment or decrement indication.

As an example of the "8 of 10" objective, suppose a pointer processor receives a pointer word that has (due to transmission errors) three of the I-bits and two of the D-bits inverted from the previous pointer value. In that case, six of the 10 I- and D-bits would be correct for an increment operation (i.e., the three inverted I-bits plus the three noninverted D-bits), while four of the 10 would be correct for a decrement operation (i.e., the two noninverted I-bits plus the two inverted D-bits). If the pointer processor met the "8 of 10" objective, it would not interpret the pointer to be indicating either an increment or decrement operation.

10. If a majority of the I-bits and a majority of the D-bits are detected to be inverted in the same pointer word, there are several ways that a pointer processor that does not meet the "8 of 10" objective can react. For example, this requirement could be literally interpreted to mean that the pointer processor must consider the pointer value to be both incremented and decremented (i.e., it would consider the contents of the H3 byte to be part of the SPE, but would ignore the byte following H3). However, the intent of the requirement is that the pointer processor not consider the pointer value to be either incremented or decremented (which is consistent with the "8 of 10" objective).

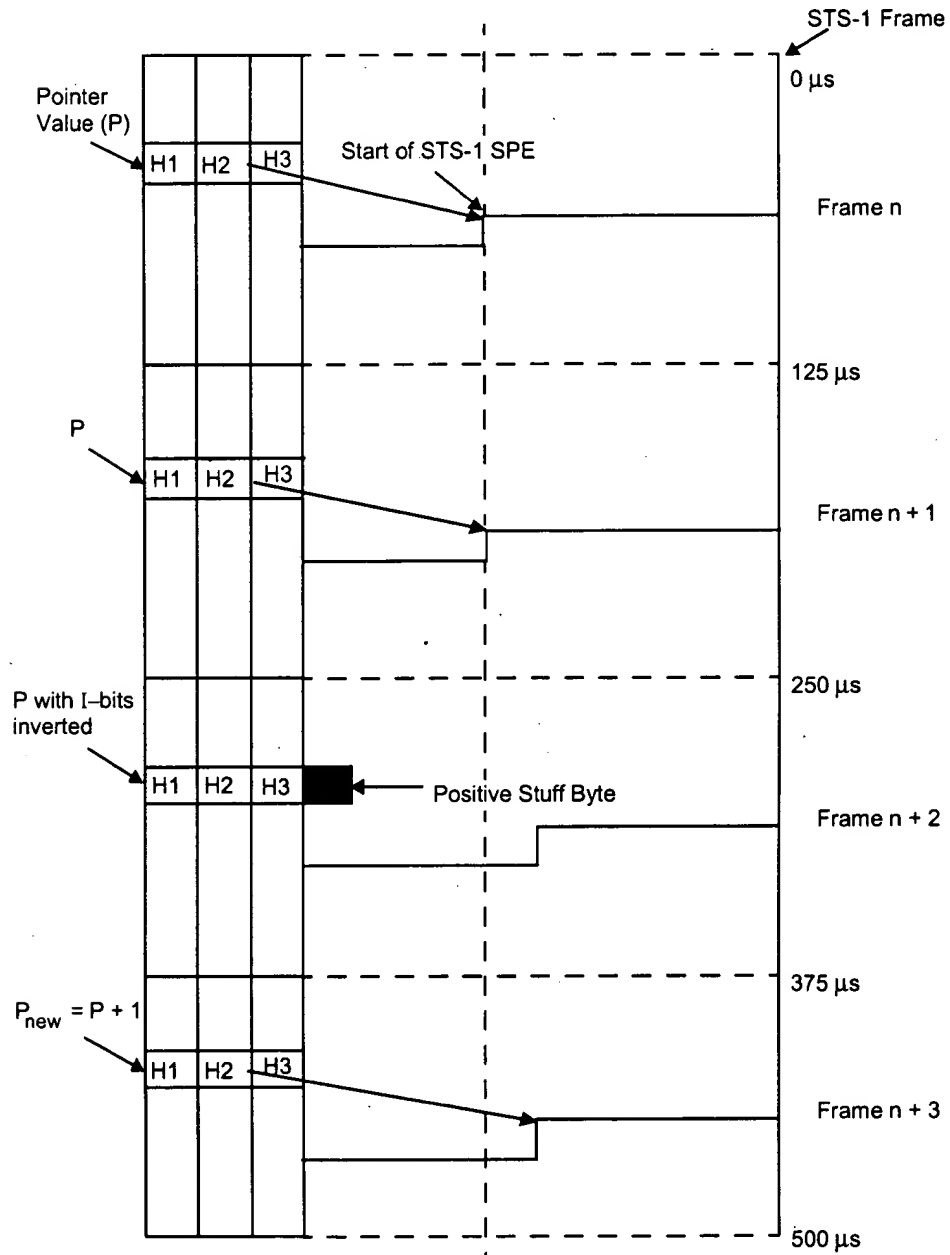


Figure 3-39. Positive STS Pointer Adjustment Operation (Increment)

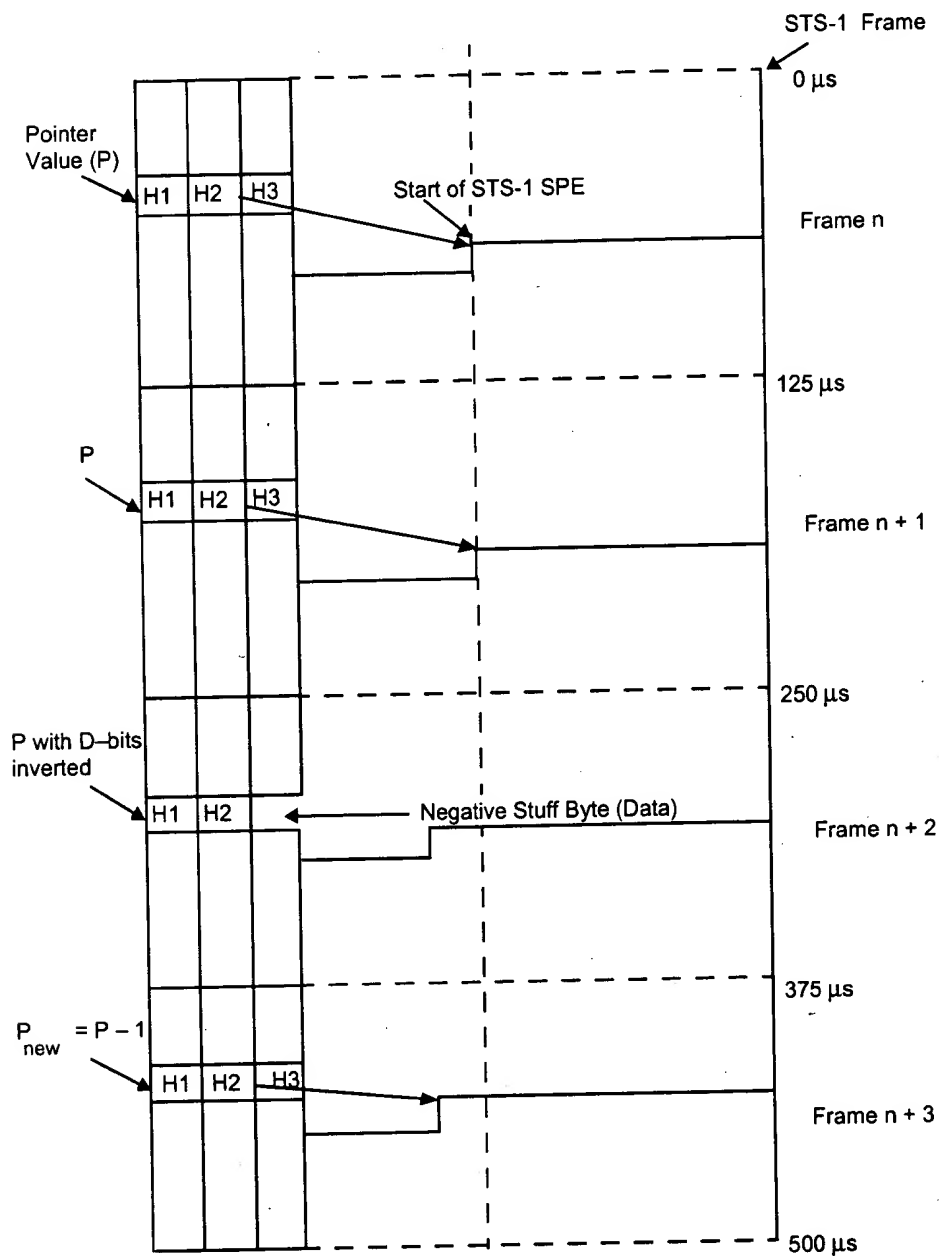


Figure 3-40. Negative STS Pointer Adjustment Operation (Decrement)

3.5.1.3 New Data Flag (NDF)

Bits 1 through 4 of the pointer word (i.e., the N-bits) carry an NDF, which can either be "normal" or "set", and which allows an arbitrary change of the pointer value due to a change in the payload.

- R3-95** [91] A normal NDF shall be indicated (during normal operation) by a '0110' code in the N-bits (see Figure 3-37). The NDF shall be set by inverting the N-bits to '1001.' The new alignment of the STS SPE shall be indicated by the pointer value accompanying the set NDF and takes effect at the offset indicated.
- R3-96** [92] The decoding at the pointer processor shall be performed by majority voting (i.e., the NDF shall be detected as being set if three or four of the N-bits match the '1001' code). If a set NDF is detected, then the coincident pointer value shall replace the current value at the offset indicated by the new pointer value.

3.5.1.4 Concatenation Indicator

Concatenation indicators contained in the Payload Pointers of the second through Nth STS-1s in an STS-Nc are used to show that those STS-1s each contain part of the STS-Nc SPE (rather than individual STS-1 SPEs).

- R3-97** [93] The first STS-1 within an STS-Nc shall have a normal pointer word.
- R3-98** [94] All subsequent STS-1s within the STS-Nc shall have their pointer values (i.e., bits 7 through 16) set to all-ones, and their N-bits set to '1001' (i.e., set NDFs).

This value of the pointer word (i.e., '1001XX11 11111111') is the concatenation indicator, and does not indicate an offset.

- R3-99** [95] A pointer processor in an NE that is transmitting or receiving an STS-Nc SPE shall perform the operations indicated by the pointer in the first STS-1 of the STS-Nc on all N of the STS-1s in that STS-Nc.

This means that if a decrement operation is transmitted or detected on the pointer in the first STS-1 of the STS-Nc, then the H3 bytes in all N of the STS-1s in the STS-Nc are considered to be negative stuff bytes. Similarly, if an increment operation is transmitted or detected, then the first N bytes of the STS-Nc envelope capacity following the last H3 byte are considered positive stuff bytes.

3.5.1.5 STS Payload Pointer Generation Rules

The pointer generation criteria from Sections 3.5.1.1 through 3.5.1.4 are summarized in this section as a set of pointer generation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

- R3-100** [96v2] The STS Payload Pointer shall be generated according to these rules:
- ... 1. During normal operation, a normal NDF is sent (i.e., the N-bits are set to '0110'), and the pointer value locates the start of the STS SPE within the STS Envelope Capacity.
 - ... 2. The pointer value shall only be changed by the operations in rules 4, 5, or 6.
 - ... 3. If an STS-Nc SPE is being transmitted, a normal pointer word is generated for the first STS-1 only. The concatenation indicator is generated in the other pointers. All operations indicated by the pointer in the first STS-1 apply to each STS-1 in the STS-Nc.
 - ... 4. If a positive stuff is needed, the current pointer value is sent with the I-bits inverted, and the subsequent positive stuff opportunity is considered an undefined byte. Subsequent pointers contain the previous pointer value incremented by one.
 - ... 5. If a negative stuff is needed, the current pointer value is sent with the D-bits inverted, and the subsequent negative stuff opportunity is overwritten with an SPE byte. Subsequent pointers contain the previous pointer value decremented by one.
 - ... 6. If the alignment of the SPE changes for any reason other than rules 4 or 5, the new pointer value shall be sent accompanied by a set NDF. The set NDF only appears in the first frame that contains the new value. The new SPE begins at the first occurrence of the offset indicated by the new pointer value.
 - ... 7. No increment or decrement operation shall be performed for three frames following any of the operations in rules 4, 5, and 6.¹¹
 - ... 8. For a nonterminated path, an incoming all-ones pointer word shall be regenerated or relayed with no more than a three-frame delay. When a non-all-ones pointer word is subsequently received, the downstream pointer shall be generated based on the pointer generation and

11. Note that in the pointer interpretation rules in the following section, there is no rule or requirement equivalent to pointer generation rule 7. If a pointer processor detects an increment or decrement operation within three frames after another pointer change operation (e.g., due to transmission errors), it can either ignore that operation or interpret it as a valid operation.

interpretation criteria summarized in this requirement (**R3-100** [96v2]) and **R3-102** [97].¹²

Rule 8 is applicable to LTE that processes STS pointers. If all of the LTE in a long chain of SONET NEs takes the full three frames allowed in Rule 8, then the detection and termination of an AIS-P defect at the STS PTE could be significantly delayed. Therefore the following objective is applicable.

O3-101 [908] LTE that processes STS pointers should regenerate or relay an incoming all-ones pointer word with no more than a one-frame delay.

Note that LTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, and therefore the incoming STS SPE associated with the all-ones pointers cannot be expected to be passed through unaltered (see Section 6.2.1.2.2).

3.5.1.6 STS Payload Pointer Interpretation

The pointer interpretation criteria from Sections 3.5.1.1 through 3.5.1.4 are summarized in this section as a set of pointer interpretation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

R3-102 [97] The STS Payload Pointer shall be interpreted according to these rules:

- ... 1. During normal operation, the pointer value locates the start of the STS SPE within the STS Envelope Capacity.
- ... 2. Any variation from the current pointer value shall be ignored unless a consistent new value is received three times consecutively, or the variation is one of the operations in rules 4, 5, or 6. Any consistent new value received three times in succession shall replace the current value at the offset indicated by the new pointer value.
- ... 3. If the pointer word contains the concatenation indicator, then the operations performed on that STS-1 are identical to those performed on the first STS-1 within the STS-Nc. Rules 4 and 5 do not apply to this pointer word.
- ... 4. If an increment is detected, then the byte following H3 shall be considered a positive stuff byte, and the current pointer value shall be incremented by one.
- ... 5. If a decrement is detected, then H3 shall be considered a negative stuff byte, and the current pointer value shall be decremented by one.

12. Since the STS Path is not terminated, the STS AIS detection and generation criteria in Section 6.2.1.2.2 are not the applicable criteria for determining the value of the pointer word to be sent downstream.

- ...
6. If a set NDF is detected, then the coincident pointer value replaces the current value at the offset indicated by the new pointer value.

3.5.2 VT Payload Pointer

Analogous to the STS Payload Pointer, the VT Payload Pointer provides a method of allowing flexible and dynamic alignment of the VT SPE within the VT Superframe (and therefore within the STS SPE), independent of the actual contents of the VT SPE. The VT Payload Pointer is contained in the V1 and V2 bytes, and designates the location of the byte where the VT SPE begins (i.e., the V5 byte). The V1 and V2 bytes can be viewed as one word, as shown in Figure 3-41. Bits 1 through 4 of the pointer word carry the New Data Flag, bits 5 and 6 indicate the size of the VT, and bits 7 through 16 carry the pointer value.

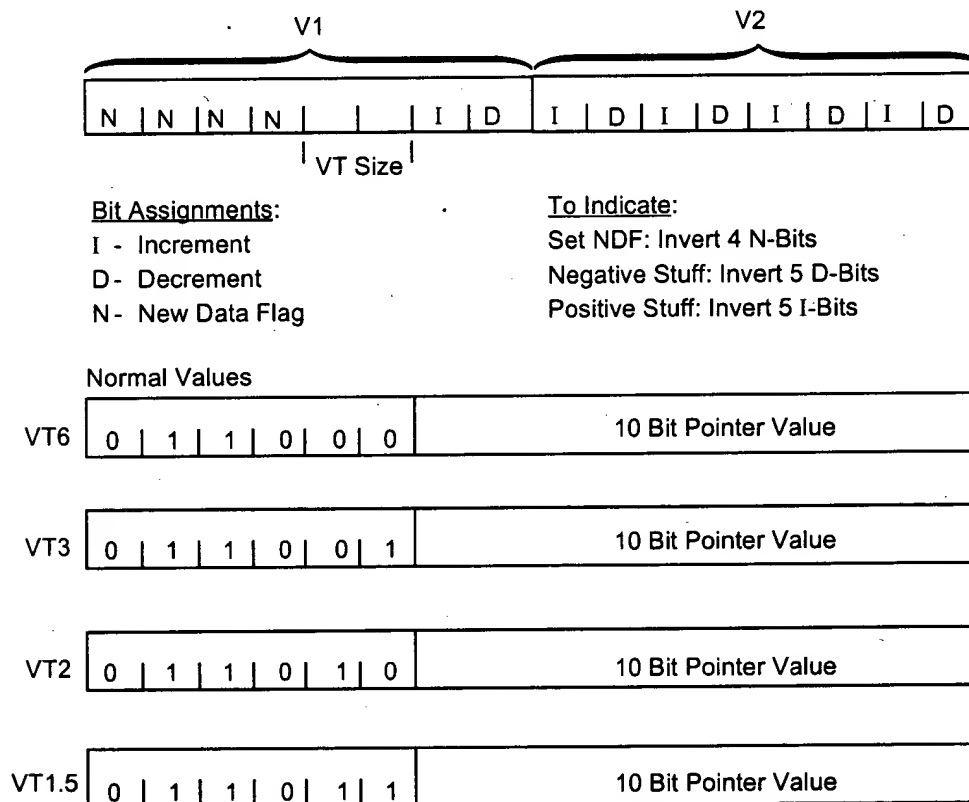


Figure 3-41. VT Payload Pointer (V1, V2) Coding

3.5.2.1 VT Pointer Value

Bits 7 through 16 of the pointer word carry the pointer value, which indicates the offset between the pointer word and the first byte of the VT SPE. The V1 through V4 bytes are not counted in the offset. The range of valid offsets is a function of the size of the VT, as shown in Figure 3-42.

- R3-103** [98] The pointer value shall be a binary number with a range of 0 to 103 (VT1.5), 0 to 139 (VT2), 0 to 211 (VT3), or 0 to 427 (VT6), and shall indicate the offset between the pointer word and the first byte of the VT SPE (as shown in Figure 3-42).

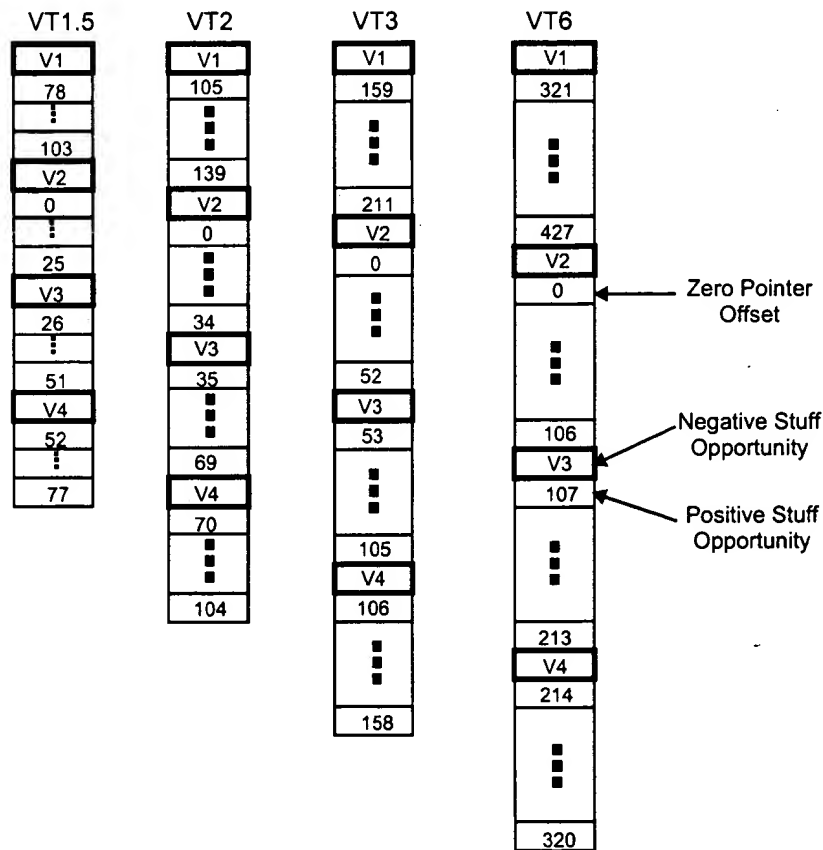


Figure 3-42. VT Pointer Offsets

3.5.2.2 VT Frequency Justification

When the frame rate of the VT SPE is greater than or less than that of the STS SPE, the alignment of the VT SPE is periodically slipped forward or back in time (using a negative or positive stuff byte) and the pointer value is adjusted by one.

- R3-104** [99] When there is a frequency offset between the frame rate of the STS SPE and that of the VT SPE, the pointer value shall be incremented or decremented as needed (although see rule 6 of **R3-113 [108v2]**), accompanied by a corresponding positive or negative stuff byte.
- R3-105** [100] A pointer increment operation shall be indicated by inverting bits 7, 9, 11, 13, and 15 (the I-bits) of the pointer word. The positive stuff byte shall appear immediately after the V3 byte in the superframe containing the inverted I-bits, as shown in Figure 3-42.

The value contained in the positive stuff byte is undefined.

- R3-106** [101] A pointer decrement operation shall be indicated by inverting bits 8, 10, 12, 14, and 16 (the D-bits) of the pointer word. The V3 byte shall be used as the negative stuff byte, (i.e., it is used to carry an SPE byte in the superframe containing the inverted D-bits), as shown in Figure 3-42.

The following criteria are applicable to the receiving VT pointer processor, and provide protection against errors on the I- and D-bits.

- R3-107** [102] If **O3-108 [103]** (the "8 of 10" objective) is not met, then the increment decision shall be made by a majority vote of the I-bits, and the decrement decision shall be made by a majority vote of the D-bits.
- O3-108** [103] The increment/decrement decision should be made at the receiver by a match of 8 or more of the 10 I- and D-bits to either the increment or decrement indication.

3.5.2.3 VT Size Indicator

Bits 5 and 6 of the VT pointer word are used to indicate the size of the VT.

- R3-109** [104] Bits 5 and 6 of the VT Payload Pointer shall indicate the size of the VT using the code '00' (VT6), '01' (VT3), '10' (VT2), or '11' (VT1.5).

The VT size codes are also listed in Table 3-5, along with the VT pointer value ranges applicable for each size of VT.

Table 3-5. VT Size Indicator

Size Bits 5 and 6	Designation	VT Pointer Range
00	VT6	0 through 427
01	VT3	0 through 211
10	VT2	0 through 139
11	VT1.5	0 through 103

3.5.2.4 New Data Flag (NDF)

Bits 1 through 4 of the pointer word carry the NDF. The NDF allows an arbitrary change of the pointer value, and also allows an arbitrary change of the size of the VTs in a VT group (i.e., due to a change in the payload). If there is a change in the size of one VT in a VT group, then there is implicitly a simultaneous change in the size of all of the VTs in that group.

- R3-110** [105] A normal NDF shall be indicated (during normal operation) by a '0110' code in the N-bits (see Figure 3-41). The NDF shall be set by inverting the N-bits to '1001.' The new alignment of the VT SPE shall be indicated by the pointer value accompanying the set NDF and takes effect at the offset indicated.
- R3-111** [106] The decoding at the pointer processor shall be performed by majority voting (i.e., the NDF shall be detected as being set if three or four of the N-bits match the '1001' code). If a set NDF is detected, then the coincident pointer value shall replace the current value at the offset indicated by the new pointer value.
- R3-112** [107] If a new size of VT is transmitted, then all 1 to 4 (depending on the new size) of the VT Payload Pointers in the VT group shall simultaneously indicate a set NDF and the same new size. The new size shall take effect immediately.¹³

13. Note that when VT AIS is being transmitted, the entire VT Superframe contains all-ones, including the N-bits and size bits in the VT pointer word. Since the NDF and size bits are needed in all of the VTs of the new size to indicate to the receiving pointer processor that a new size of VTs is being transmitted, this requirement implies that VT AIS cannot be transmitted in the first VT Superframes containing the new VTs.

3.5.2.5 VT Payload Pointer Generation Rules

The pointer generation criteria from Sections 3.5.2.1 through 3.5.2.4 are summarized in this section as a set of pointer generation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

R3-113 [108v2] The VT Payload Pointer shall be generated according to these rules:

- ... 1. During normal operation, a normal NDF is sent (i.e., the N-bits are set to '0110'), the size bits indicate the size of the VT, and the pointer value locates the start of the VT SPE within the VT Envelope Capacity.
- ... 2. The pointer value shall only be changed by the operations in rules 3, 4, or 5.
- ... 3. If a positive stuff is needed, the current pointer value is sent with the I-bits inverted, and the subsequent positive stuff opportunity is considered an undefined byte. Subsequent pointers contain the previous pointer value incremented by one.
- ... 4. If a negative stuff is needed, the current pointer value is sent with the D-bits inverted, and the subsequent negative stuff opportunity is overwritten with an SPE byte. Subsequent pointers contain the previous pointer value decremented by one.
- ... 5. If the alignment of the SPE changes for any reason other than rules 3 or 4, the new pointer value shall be sent accompanied by a set NDF. The set NDF only appears in the first superframe that contains the new value. The new SPE begins at the first occurrence of the offset indicated by the new pointer value.
- ... 6. No increment or decrement operation shall be performed for three superframes following any of the operations in rules 3, 4, and 5.¹⁴
- ... 7. For a nonterminated path, an incoming all-ones pointer word shall be regenerated or relayed with no more than a three-superframe delay. When a non-all-ones pointer word is subsequently received, the downstream pointer shall be generated based on the pointer generation and interpretation criteria summarized in this requirement (R3-113 [108v2]) and R3-115 [109].¹⁵

14. Note that in the pointer interpretation rules in the following section, there is no rule or requirement equivalent to pointer generation rule 6. If a pointer processor detects an increment or decrement operation within three superframes after another pointer change operation (e.g., due to transmission errors), it can either ignore that operation or interpret it as a valid operation.

15. Since the VT Path is not terminated, the VT AIS detection and generation criteria in Section 6.2.1.2.3 are not the applicable criteria for determining the value of the pointer word to be sent downstream.

- ...
8. If the size of the VTs within a VT group is to change, then the NDFs in all of the VTs of the new size (in that VT group) are set simultaneously.

Rule 7 is applicable to STS PTE that processes VT pointers. If all of the STS PTE in a long chain of SONET NEs takes the full three superframes allowed in Rule 7, then the detection and termination of an AIS-V defect at the VT PTE could be significantly delayed. Therefore the following objective is applicable.

- O3-114** [909] STS PTE that processes VT pointers should regenerate or relay an incoming all-ones pointer word with no more than a one-superframe delay.

Note that STS PTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, and therefore the incoming VT SPE associated with the all-ones pointers cannot be expected to be passed through unaltered (see Section 6.2.1.2.3).

3.5.2.6 VT Payload Pointer Interpretation Rules

The pointer interpretation criteria from Sections 3.5.2.1 through 3.5.2.4 are summarized in this section as a set of pointer interpretation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

- R3-115** [109] The VT Payload Pointer shall be interpreted according to these rules:

- ...
1. During normal operation, the pointer value locates the start of the VT SPE within the VT Envelope Capacity.
 - ...
 2. Any variation from the current pointer value shall be ignored unless a consistent new value is received three times consecutively, or the variation is one of the operations in rules 3, 4, or 5. Any consistent new value received three times in succession shall replace the current value at the offset indicated by the new pointer value.
 - ...
 3. If an increment is detected, then the byte following V3 shall be considered a positive stuff byte, and the current pointer value shall be incremented by one.
 - ...
 4. If a decrement is detected, then the V3 byte shall be considered a negative stuff byte, and the current pointer value shall be decremented by one.
 - ...
 5. If a set NDF is detected, then the coincident pointer value replaces the current value at the offset indicated by the new pointer value.
 - ...
 6. If the equipment has the capability to correctly process different VT sizes based on the received VT size bits, and a set NDF and an

arbitrary new size of VT are received simultaneously in all of the VTs within a VT group, then the coincident pointer values and sizes shall replace the current pointer values and sizes at the offsets indicated in the new pointers.

- ...
7. If the equipment has the capability to correctly process different VT sizes based on the received VT size bits, then any variation from the current VT size shall be ignored unless consistent valid pointers indicative of a new VT size are received three times consecutively in all of the (new) VTs within a VT group, or the variation is the operation in rule 6. The VT size associated with such pointers received three times in succession shall replace the current size immediately.

As indicated above, the applicability of VT pointer interpretation rules 6 and 7 depend on the functionality of the receiving equipment. Some equipment may be capable of correctly processing different VT sizes based on the received VT size bits. For example, STS PTE that processes the incoming VT pointers and passes VTs (in groups) to other STS PTE within the same NE for transmission on an outgoing SONET signal could detect changes in the VT size bits and change its pointer processing functions accordingly. If such a capability is provided, then rules 6 and 7 would be applicable. Other equipment may not be capable of processing different size VTs correctly using the VT size bits. For example, a SONET ADM that is only equipped to support DS1 and DS3 low-speed interfaces would only be expected to be capable of processing VT1.5s. If that ADM received VTs of any other size, it would be expected to declare LOP-V alarms for all of the affected VT1.5s. In such a case, the LOP defect detection criteria in Section 6.2.1.1.3 would be the applicable criteria, and rules 6 and 7 would not apply.

In addition, the phrase "valid pointers indicative of a new VT size" in pointer interpretation rule 7 can be interpreted various ways, and is intended to allow for any existing designs while encouraging the development of new designs that can detect incoming VT size changes that are not detected via rule 6 (e.g., due to transmission errors). The intended interpretation is that the pointer words in at least one of the VTs of the new size should contain normal NDFs, consistent and valid size bits, and a constant in-range pointer value (i.e., a normal pointer word) for three consecutive VT Superframes. The remaining VTs could have either consistent normal pointer words or consistent all-ones pointer words (i.e., AIS) in their corresponding VT Superframes.

4. Physical Layer

This section describes optical and electrical Physical Layer parameters that enable multisupplier compatibility for SONET signals at the Physical Layer. Section 4.1 provides an overview of the Physical layer classifications. Section 4.2 contains the optical parameter requirements for the application categories defined in Section 4.1. Section 4.3 describes a methodology that may be used for engineering a single-mode fiber optic system, and Section 4.4 contains the SONET electrical interface specifications. Interface parameters for SONET transport systems utilizing fiber amplifiers are for further study.

4.1 Physical Layer Classifications

SONET signals may be transported by either electrical or optical means, and SONET NEs may have either electrical or optical interfaces (or both). While technical considerations limit the feasibility of electrical transport to short distances and relatively low bit rates, the flexibility of optical transmission allows a wide range of possible applications and corresponding implementations. However, to simplify the development of compatible multisupplier SONET optical systems, it is desirable to define a relatively small set of application categories and corresponding sets of optical interface specifications. This allows the ability to obtain some degree of optical device commonality among the various applications and to balance economic and technical considerations in developing compatible multisupplier SONET equipment.

In this document, the following broad application categories are used:

- Long Reach (LR) optical interfaces refer to optical sections with system loss budgets from 10 dB to 28 dB at OC-1 and OC-3, and to 24 dB at OC-12 and OC-48. Typical of long-haul telecommunications systems, LR interfaces are based on high-power (e.g., 500 μ W or -3 dBm) Multi-Longitudinal Mode (MLM) or Single-Longitudinal Mode (SLM) lasers.
- Intermediate Reach (IR) optical interfaces refer to optical sections with system loss budgets from 0 dB to 12 dB. Typically, low-power (e.g., 50 μ W or -13 dBm) SLM or MLM laser transmitters are used.
- Short Reach (SR) optical interfaces refer to optical sections having system loss budgets from 0 dB to 7 dB. Depending on the SONET hierarchical level, SR transmitters may be either Light-Emitting Diodes (LEDs) or low-power MLM lasers.
- Electrical interfaces, compatible with DSX-type interfaces for the asynchronous hierarchy, are defined for the STS-1 and STS-3 levels.

Optical parameters are defined for most levels of the SONET hierarchy (all levels with the exception of OC-24) in the three broad application categories (LR, IR, and SR). Within each of the application categories, it is possible to consider the use of either dispersion-unshifted single-mode fiber (EIA/TIA Class IVa) or dispersion-shifted single-mode fiber

(EIA/TIA Class IVb). GR-20-CORE, *Generic Requirements for Optical Fiber and Fiber Optic Cable*, contains generic requirements for optical fiber and optical fiber cables. This GR contains requirements based on ITU-T Recommendation G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*, for LR, IR, and SR optical systems operating with nominal 1310-nm sources on dispersion-unshifted single-mode fiber and for LR and IR optical systems operating with nominal 1550-nm sources on both dispersion-unshifted and dispersion-shifted single-mode fiber. These applications are considered to be the predominant current and near-term future applications for SONET systems in the BCC networks. To categorize the applications described above according to nominal source wavelength and nominal fiber zero-dispersion wavelength, it is useful to define a classification scheme as shown in Table 4-1.¹

Table 4-1. Application Categories by Nominal Spectral Attributes

Application	Nominal Source Wavelength (nm)	Nominal Fiber Zero-Dispersion Wavelength (nm)
LR-1	1310	1310*
LR-2	1550	1310*
LR-3	1550	1550**
IR-1	1310	1310*
IR-2	1550	1310*
SR	1310	1310*

* Dispersion-unshifted single-mode fiber per GR-20-CORE

** Dispersion-shifted single-mode fiber per GR-20-CORE

4.2 Optical Parameter Definitions and Interface Requirements

4.2.1 General

To specify SONET optical parameters for the various application categories, optical system interfaces can be represented as shown in Figure 4-1. Point S is a reference point on the

1. Note that two Short Reach application categories are expected to be defined in ANSI T1.105.06, *SONET Physical Layer Specifications*. These are the SR-0 category, which uses a nominal 1310 nm source wavelength and multimode fiber, and the SR-1 category which corresponds to the SR category in this document. In addition, ANSI is also expected to define optical parameters for the OC-24 and Optical VT Group (OVTG) rates, and electrical parameters for the VT1.5 rate. The need to include some or all of these additional parameters in future issues of this document is under study.

optical fiber just after the transmitter (Tx) optical connector (C_{Tx}). Point R is a reference point on the optical fiber just before the receiver (Rx) optical connector (C_{Rx}). Points S and R provide a convenient separation of the optical link into a transmitter subsection, a receiver subsection, and an optical connection subsection. Optical parameters are specified for the transmitter at point S, for the receiver at point R, and for the optical connection between points S and R. Any additional connectors within the optical link, such as those at a fiber distribution frame, are considered part of the fiber link and are associated with the optical connection.

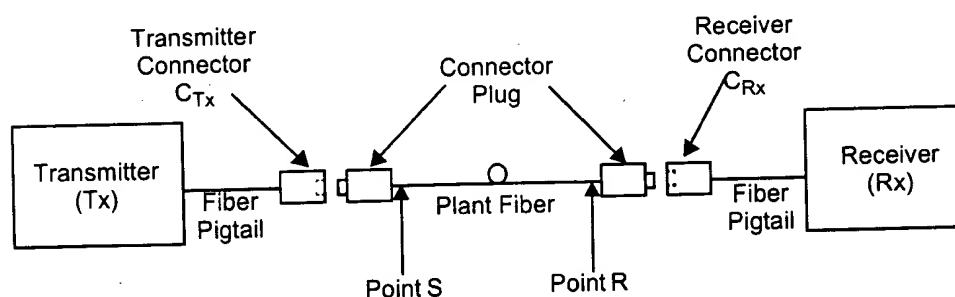


Figure 4-1. Optical System Interfaces (Points S and R)

All parameter values specified are worst-case, end-of-life values and are to be met over the ranges of operating conditions that Section 7.1.1 describes. The parameters are specified relative to an optical system design objective of a BER not worse than 10^{-10} for the extreme case of optical path attenuation and dispersion conditions for each application specified. Optical parameters for systems having design objectives not worse than 10^{-n} , where $n > 10$, are for further study. Separate equipment margins are not specified, and it is assumed that transmitters, receivers, and the optical cable plant individually meet the specifications that follow in Tables 4-3 through 4-11. The intent of the interface specifications in these tables is to enable multisupplier compatibility. In some cases, use of these specifications may lead to more conservative optical section designs than could be obtained through supplier-proprietary interfaces, the use of statistical design approaches for engineering an optical link, or in environments more constrained than those that Section 7.1.1 describes. Further, for all OC-48 applications, multisupplier compatibility cannot be guaranteed on the basis of the current specifications, which may not adequately account for laser dynamic effects. These effects are for further study. Tables 4-3 through 4-11 indicate optical parameters that are not specified or that are not appropriate with the notation "NA" (Not Applicable).

4.2.2 Optical Line Coding

- R4-1 [110] For all SONET optical system interfaces described, binary Non-Return-to-Zero (NRZ) optical line coding shall be used.

4.2.3 Reflections

Refractive index discontinuities along the optical path may cause reflections. If not controlled, reflections can degrade system performance through their disturbing effect on the operation of the laser or through multiple reflections that lead to interferometric noise at the receiver. While the performance degradation from single-point reflections on the laser can be reduced by optical isolation of the transmitter, multiple reflection effects are much more difficult to predict and control because they depend in a complicated fashion on the number, spacing, and reflectance value (R) of the discrete reflectors as well as on the optical source type and its spectral characteristics. In general, reflection-induced degradation increases with system bit-rate, optical source coherence, and fiber dispersion. Minimizing individual reflectance reduces the impact of single and multiple reflections on system performance. By enforcing reflectance requirements on individual components placed in the fiber optic transmission span, and by requiring system performance to have a tolerance to specified reflectance values from those components, the effects of fiber optic system reflection noise can be minimized.

- R4-2 [111] To ensure the capability for upgrading SONET transport systems to high bit-rates, all splices (see GR-765-CORE, *Generic Requirements for Single Fiber Single-Mode Optical Splices and Splicing Systems*), connectors (see GR-326-CORE, *Generic Requirements for Single-Mode Optical Fiber Connectors*), attenuators (see GR-910-CORE, *Generic Requirements for Fiber Optic Attenuators*), couplers, and Wavelength-Division-Multiplexing (WDM) components (see GR-1209-CORE, *Generic Requirements for Fiber Optic Branching Components*) intended for installation in new facilities shall meet the reflectance requirements specified in the referenced documents.

In addition, to accommodate embedded facilities that may have reflection performance worse than the current recommendations, and to recognize the greater tolerance of low bit-rate systems to reflections, it is necessary that SONET system transmitters and receivers tolerate higher levels of reflections than would be generated by the discrete reflections described above. Specific requirements relate to discrete reflections along the optical path, to values of the system Optical Return Loss (ORL), which accounts for both discrete reflections and distributed reflections along the fiber itself, and to reflections from the receiver. Values, based on the specifications in ITU-T G.957, are given in Tables 4-3 through 4-11. For embedded facilities that may have components with reflectance worse than these specifications, it is an individual BCC decision to determine the necessity of

replacing these components or reducing their number to upgrade each facility to the highest anticipated line rate and desired application.

The receiver reflectance value of -27 dB that some applications in Tables 4-3 through 4-11 specify is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations involving things such as multiple connectors and splices. Systems employing fewer or higher-performance optical components produce fewer multiple reflections and, consequently, are able to tolerate receivers having higher reflectance values. As an extreme example, if only two connectors exist in the system, a -14 dB receiver reflectance is considered acceptable. Therefore, at the option of each BCC and depending on the specific application intended, maximum receiver reflectance values lower (more negative) than -14 dB in Tables 4-3 through 4-11 are required as follows. Note that no receiver reflectance criteria exist for applications where "NA" is shown in Tables 4-3 through 4-11.

- R4-3** [112] The receiver reflectance shall be less than (more negative than) the value is listed under "Max. Receiver Reflectance" in Tables 4-3 through 4-11.

System ORL is the ratio (in dB) of the optical power (P_1) arriving downstream at a system interface to the optical power (P_2) reflected back upstream to the same interface. This includes the reflected power contributions from all system components downstream from the interface.

$$\text{System ORL} = 10 \log_{10} \left[\frac{P_1}{P_2} \right] \text{ dB.} \quad (4-1)$$

To measure reflectance between points S and R and ORL at points S and R, the measurements are taken at the end-faces of the connector plugs of the plant fiber that Figure 4-1 shows. Since such measurements do not include the performance of the respective connectors in the operational system, it may be assumed that the connectors C_{Tx} and C_{Rx} have nominal reflectance values for the specific connector type used.

- R4-4** [113] SONET optical transmitters and receivers shall operate properly in the presence of the worst-case combination of discrete reflectance, including receiver reflectance, system ORL, and minimum optical path attenuation values given in Tables 4-3 through 4-11. Proper system operation results in a system power penalty less than 1 dB under worst-case reflection conditions.

For some applications in Tables 4-3 through 4-11, reflections are not expected to limit system performance. For these applications, the minimum ORL (ORL_{min}) entries are indicated by "NA".

- O4-5** [114] For all applications in Tables 4-3 through 4-11, the optical system should operate properly in the presence of a -8.5 dB reflection (i.e., maintain a system power penalty less than 1 dB).

4.2.4 Transmitter

This document specifies transmitter requirements that include the range of operating wavelengths, spectral width, range of coupled transmit power, and extinction and side-mode suppression ratios. In addition, transmitter pulse shapes are specified by the mask of the eye diagram at point S.

Tables 4-3 through 4-11 list required transmitter characteristics for the various SONET applications. Depending on the attenuation/dispersion characteristics and hierarchical level of each application in Tables 4-3 through 4-11, feasible transmitter devices include LEDs, MLM lasers, or SLM lasers. This document indicates one or more nominal device types for each application. This indication is not intended to be a requirement on the source type, and any device meeting the transmitter characteristics specified may be substituted for the nominal device type without degradation in system performance.

- R4-6** [115] The transmitter central wavelength and coupled transmit power shall be within the appropriate ranges listed in Tables 4-3 through 4-11.
- R4-7** [116] The spectral width of the transmitter shall be less than or equal to the appropriate value listed in Tables 4-3 through 4-11.
- R4-8** [117] The side-mode suppression ratio and extinction ratio of the transmitter shall be greater than or equal to the appropriate values listed in Tables 4-3 through 4-11.

For some applications using MLM laser transmitters (e.g., IR-1 of Table 4-9), two possible spectral width ranges and corresponding central wavelength ranges are given. This permits the use of MLM transmitters having wider spectral widths when the transmitter operating wavelength range is closer to the fiber zero-dispersion wavelength for the fiber type specified for the application. However, in dispersion-limited systems, transmitters having narrower spectral widths correspond to larger values of maximum allowed system dispersion (ps/nm) and, hence, to longer optical spans. The ranges of central wavelengths that the tables specify are consistent with the values given in ITU-T G.957.

- O4-9** [118] In Tables 4-4, 4-5, and 4-9, it is an objective for MLM transmitters to meet the narrower spectral width specifications for those applications that list two possible values for $\Delta\lambda_{rms}$.

The possible need for additional or alternative specifications to account for dynamic laser effects (e.g., laser chirp) and other spectral characteristics of SLM devices is under study.

4.2.4.1 Spectral Characteristics

The following definitions apply to the parameter values that Tables 4-3 through 4-11 specify.

λ_{Tnom} – The nominal value of the central operating wavelength. It is the wavelength where the effective optical power resides (see FOTP-127, *Spectral Characterization of Multimode Laser Diodes*).

λ_{Tmin} , λ_{Tmax} – The range of central operating wavelengths. It is defined as the allowable range of transmitter central wavelengths, around λ_{Tnom} , under worst-case variations due to manufacturing, temperature, aging, and reflections. It is measured under fully modulated conditions and in the operating environment that Section 7.1.1 describes.

$\Delta\lambda_{rms}$ – Root-mean-square (rms) spectral width. For LEDs and MLM lasers, spectral width is specified by $\Delta\lambda_{rms}$ under the operating conditions that Section 4.2.1 defines and in the presence of the worst-case reflections that Tables 4-3 through 4-11 specify for each SONET application (see FOTP-127). The measurement of rms spectral width accounts for modes out to and including those 20 dB down from the peak mode. $\Delta\lambda_{rms}$ does not apply to SLM transmitters.

$\Delta\lambda_{20}$ – Full spectral width measured 20 dB down from the maximum of the central wavelength peak of an SLM transmitter operating under fully modulated conditions in the presence of the worst-case reflections. $\Delta\lambda_{20}$ does not apply to MLM lasers or to LEDs.

SSR_{min} – Minimum acceptable value of the Side-mode Suppression Ratio (SSR). SSR is defined as the ratio (in dB) of the average optical power in the dominant longitudinal mode (M1) of an SLM laser to the optical power in the most significant side-mode (M2) under fully modulated conditions, in the presence of worst-case reflections, as follows:

$$SSR = 10\log_{10}\left(\frac{M1}{M2}\right)\text{dB.} \quad (4-2)$$

4.2.4.2 Coupled Transmit Power

P_{Tmax} , P_{Tmin} – The maximum and minimum of the allowed range of average coupled transmitter power, P_T , at point S for a pseudo-random data sequence coupled into the fiber from the optical transmitter. These are worst-case values to account for manufacturing variances, drifts due to temperature variations and aging effects, and operation with the specified minimum value of the extinction ratio (see OFSTP-2, *Effective Transmitter Output Power Coupled into Single-Mode Fiber Optic Cable*).

4.2.4.3 Extinction Ratio

r_{emin} – Minimum value (min) of the extinction ratio r_e . It is the minimum acceptable value of the ratio (in dB) of the average optical energy in a logic one level $E_R(1)$ to the average optical energy in a logic zero level $E_R(0)$ measured under fully modulated conditions in the presence of worst-case reflections, as follows:

$$r_e = 10 \log_{10} \left[\frac{E_R(1)}{E_R(0)} \right] \text{ dB.} \quad (4-3)$$

4.2.4.4 Mask of the Eye Diagram

X1, X2, Y1 – Parameters specifying the (normalized) mask of the transmitter eye diagram. See Figures 4-2 and 4-3.

Values for these parameters are determined with respect to a low-pass reference filter that need not represent the noise filter of the transmission equipment optical receiver. An acceptable reference filter for qualifying transmitted pulses through an eye diagram measurement is a fourth-order Bessel-Thomson transfer function given by:

$$H(p) = (1/105) [105 + 105y + 45y^2 + 10y^3 + y^4] \quad (4-4)$$

where,

$$p = j\omega / \omega_r, \quad y = 2.1140p, \quad \omega_r = 2\pi f_r, \quad (4-5)$$

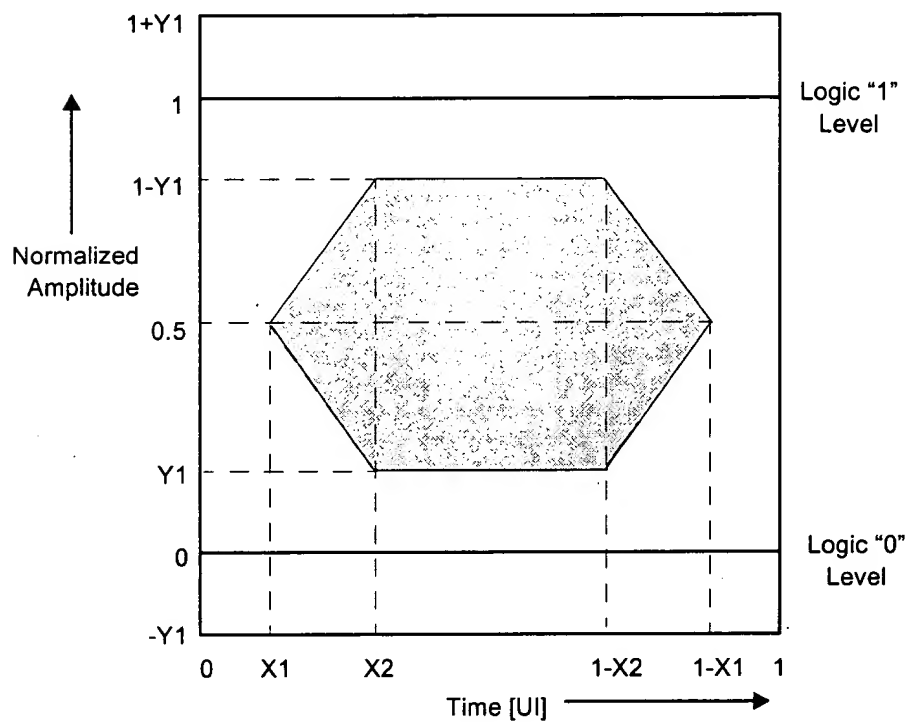
$$f_r = 0.75 f_0, \quad f_0 = \text{bit rate.} \quad (4-6)$$

For this filter, the nominal attenuation at the reference frequency, f_r , is 3 dB. The corresponding attenuation and group delay distortion at various frequencies are given in Table 4-2. The allowable deviation of the nominal attenuation values to account for manufacturing tolerances of the optical reference receiver components as well as details regarding the measurement setup for the eye diagram are under study in ITU-T Study Group 15.

- R4-10** [119] Transmit pulses, referenced to a noise filter with transfer characteristics as given in Equation 4-4, shall fall within the mask as defined in Figures 4-2 and 4-3.

Table 4-2. Attenuation and Group Delay Distortion as a Function of Frequency

f/f_0	f/f_r	Attenuation (dB)	Group Delay Distortion (UI)
0.15	0.2	0.1	0
0.3	0.4	0.4	0
0.45	0.6	1.0	0
0.6	0.8	1.9	0.002
0.75	1.0	3.0	0.008
0.9	1.2	4.5	0.025
1.0	1.33	5.7	0.044
1.05	1.4	6.4	0.055
1.2	1.6	8.5	0.10
1.35	1.8	10.9	0.14
1.5	2.0	13.4	0.19
2.0	2.67	21.5	0.30



Rates	X1	X2	Y1
OC-1 and OC-3	0.15	0.35	0.20
OC-12	0.25	0.40	0.20

Figure 4-2. SONET Eye Diagram Mask (OC-1 to OC-12)

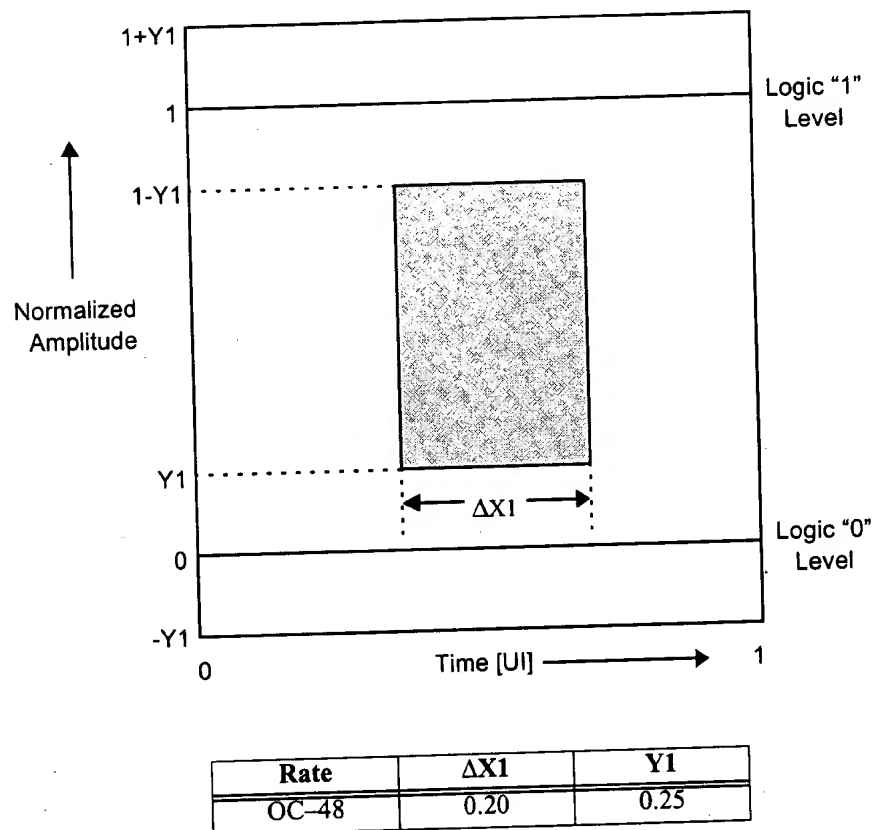


Figure 4-3. SONET Eye Diagram Mask (OC-48)

4.2.5 Receiver

Receiver characteristics specified in this document are receiver sensitivity, receiver overload, receiver reflectance, and the allowable power penalty from the combined effects of dispersion, reflections, and jitter.

P_{Rmin} , P_{Rmax} – Minimum and maximum values of the average received power, P_R in dBm, at point R to achieve a 10^{-10} BER. These values take into account power penalties caused by use of a transmitter with worst-case transmitter spectral width, extinction ratio, pulse shape characteristics, and operating wavelength range specified for each application described, and they include the effects of drifts due to temperature variations and aging.

P_O – Optical path power penalty (in dB). P_O accounts for the total degradation along the optical path between points S and R (from reflections, jitter, intersymbol interference, mode-partition noise, and laser chirp).

R4-11 [120] The minimum acceptable value for receiver sensitivity shall equal the values P_{Rmin} specified in Tables 4-3 through 4-11.

R4-12 [121] The receiver shall accommodate an optical path power penalty of at least P_O for each application specified.

In actual implementations, it is not expected that worst-case conditions for dispersion, single reflection, multiple reflection, and jitter will occur simultaneously. Therefore, for purposes of requirement verification, the power penalty shall be less than P_O under worst-case conditions for each of the aforementioned, tested separately.

O4-13 [122] It is an objective that the receiver overload point equal or exceed the value of P_{Rmax} given in Tables 4-3 through 4-11.

For most of the SONET signal rates and the application categories defined in Section 4.1, a receiver's conformance to the preceding objective permits worst-case engineering of the optical link without the use of optical attenuators for the extreme cases of minimum attenuation and maximum output power. Note however, that the use of attenuators may still be necessary in Long Reach applications at the OC-48 rate.

4.2.6 Optical Path

The attenuation ranges and dispersion parameter values in Tables 4-3 through 4-11 are expected to be sufficient to cover current and near-term future SONET installations in a typical BCC network for intra-office, inter-office short-haul (including loop feeder), and inter-office long-haul applications. To provide flexibility in implementing multisupplier SONET systems, some overlap in attenuation range is provided between the SR and IR applications and between the IR and LR applications.

The attenuation values are worst-case values including losses from splices, connectors, attenuators, or other optical devices, and any additional cable margin to cover allowances for future modifications (e.g., additional splices) to the cable plant, fiber cable performance variations due to environmental factors, and any degradation of connector or other optical device between points S and R.

Dispersion specifications are derived through consideration of the allowable pulse broadening as a fraction of the timeslot width, ϵ ,

$$\epsilon = 10^{-6} \times B \times D_{SRmax} \times \delta\lambda \quad (4-7)$$

where B is the system bit-rate (in Mb/s), D_{SRmax} is the maximum dispersion (in ps/nm) between points S and R, and $\delta\lambda$ is the spectral width (in nm). For LEDs and MLM lasers, $\delta\lambda = \Delta\lambda_{rms}$; for SLM lasers, assuming Gaussian line shapes, $\delta\lambda = \Delta\lambda_{20} / 6.07$.

For a given dispersion power penalty, ϵ has a maximum value. Specifying a value for ϵ allows calculation of D_{SRmax} that is consistent with the specified transmitter spectral width. For Intersymbol Interference (ISI) induced power penalties associated with LEDs and SLM lasers, the value $\epsilon = 0.306$ is assumed. To account for mode-partition noise in MLM lasers, the value $\epsilon = 0.115$ is assumed. These values are consistent with a total 1-dB dispersion power penalty predicted by empirical and analytic models. In the cases of the LR-2 applications for OC-48 (i.e., system operation in the high dispersion regime), the value $\epsilon = 0.491$, corresponding to a 2-dB dispersion power penalty, is assumed.

The maximum dispersion values in Tables 4-3 through 4-11 may not adequately account for the dispersion effects resulting from laser chirp. The necessity for modifying the epsilon model as well as the values of D_{SRmax} for high-bit-rate, long-haul systems to accommodate chirp-induced dispersion power penalty is under study.

In several cases in Tables 4-3 through 4-11, the system is considered to be limited by attenuation rather than by dispersion. These cases are indicated by "NA" (Not Applicable) in the D_{SRmax} entries.

Table 4-3. LR OC-1 Optical Parameters

Parameter	LR-1		LR-2	LR-3		Unit
Transmitter	MLM	SLM	SLM	MLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1280-1335	1280-1335	1480-1580	1480-1580	1480-1580	nm
$\Delta\lambda_{rms}$	13	NA	NA	5	NA	nm
$\Delta\lambda_{20}$	NA	1	1	NA	1	nm
SSR_{min}	NA	30	30	NA	30	dB
P_{Tmax}	0	0	0	0	0	dBm
P_{Tmin}	-5	-5	-5	-5	-5	dBm
r_{emin}	10	10	10	10	10	dB
Optical Path						
System ORL _{min} *	NA	NA	20	NA	NA	dB
DSR_{max}	171	NA	NA	444	NA	ps/nm
Attenuation	10-28	10-28	10-28	10-28	10-28	dB
Max Reflectance between S and R	NA	NA	-25	NA	NA	dB
Receiver						
P_{Rmax}	-10	-10	-10	-10	-10	dBm
P_{Rmin}	-34	-34	-34	-34	-34	dBm
P_O	1	1	1	1	1	dB
Max. Receiver Reflectance	NA	-25†	NA	NA	NA	dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.
- † This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.

Table 4-4. LR OC-3 Optical Parameters

Parameter	LR-1		LR-2	LR-3		Unit
Transmitter	MLM	SLM	SLM	MLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1280-1335	1280-1335	1480-1580	1534-1566**	1480-1580	nm
$\Delta\lambda_{rms}$	4	NA	NA	3 (2.5)**	NA	nm
$\Delta\lambda_{20}$	NA	1	1	NA	1	nm
SSR _{min}	NA	30	30	NA	30	dB
P _{Tmax}	0	0	0	0	0	dBm
P _{Tmin}	-5	-5	-5	-5	-5	dBm
r _{emin}	10	10	10	10	10	dB
Optical Path						
System ORL _{min} *	NA	NA	20	NA	NA	dB
DSR _{max}	185	NA	NA	246 (296)**	NA	ps/nm
Attenuation	10-28	10-28	10-28	10-28	10-28	dB
Max Reflectance between S and R	NA	NA	-25	NA	NA	dB
Receiver						
P _{Rmax}	-10	-10	-10	-10	-10	dBm
P _{Rmin}	-34	-34	-34	-34	-34	dBm
P _O	1	1	1	1	1	dB
Max. Receiver Reflectance	NA	-25†	NA	NA	NA	dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.
- ** Transmitters meeting the narrower spectral width objective are allowed a wider central wavelength range ($\lambda_{Tmin}-\lambda_{Tmax}$) of 1523-1577 nm.
- † This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.

Table 4-5. LR OC-12 Optical Parameters

Parameter	LR-1		LR-2	LR-3	Unit
Transmitter	MLM	SLM	SLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1300-1325†	1280-1335	1480-1580	1480-1580	nm
$\Delta\lambda_{rms}$	2.0 (1.7)‡	NA	NA	NA	nm
$\Delta\lambda_{20}$	NA	1	<1**	1	nm
SSR_{min}	NA	30	30	30	dB
P_{Tmax}	+2	+2	+2	+2	dBm
P_{Tmin}	-3	-3	-3	-3	dBm
r_{emin}	10	10	10	10	dB
Optical Path					
System ORL_{min} *	20	20	24	20	dB
D_{SRmax}	92 (109)‡	NA	**	NA	ps/nm
Attenuation	10-24	10-24	10-24	10-24	dB
Max. Reflectance between S and R	-25	-25	-27	-25	dB
Receiver					
P_{Rmax}	-8	-8	-8	-8	dBm
P_{Rmin}	-28	-28	-28	-28	dBm
P_O	1	1	1	1	dB
Max. Receiver Reflectance	-14	-27†	-14	-14	dB

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.
- ** Currently, there are no reliable methods for estimating chirp power penalties for SLM lasers; therefore, the values of $\Delta\lambda_{20}$ and D_{SRmax} are under study. It is expected that the value of $\Delta\lambda_{20}$ will be less than 1 nm.
- † This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.
- ‡ Transmitters meeting the narrower spectral width objective are allowed a wider central wavelength range ($\lambda_{Tmin}-\lambda_{Tmax}$) of 1296-1330 nm.

Table 4-6. LR OC-48 Optical Parameters

Parameter	LR-1	LR-2	LR-3	Unit
Transmitter	SLM	SLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1280-1335	1500-1580	1500-1580	nm
$\Delta\lambda_{rms}$	NA	NA	NA	nm
$\Delta\lambda_{20}$	1	<1**	1**	nm
SSR _{min}	30	30	30	dB
P _{Tmax}	+3	+3	+3	dBm
P _{Tmin}	-2	-2	-2	dBm
r _{emin}	8.2	8.2	8.2	dB
Optical Path				
System ORL _{min} *	24	24	24	dB
DSR _{max}	NA	1194**†	**	ps/nm
Attenuation	10-24#	10-24#	10-24#	dB
Max. Reflectance between S and R	-27	-27	-27	dB
Receiver				
P _{Rmax}	-9	-9	-9	dBm
P _{Rmin}	-27	-28	-27	dBm
P _O	1	2	1	dB
Max. Rcvr. Reflectance	-27†	-27†	-27†	dB

* For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.

** Currently, there are no reliable methods for estimating chirp power penalties for SLM lasers; therefore, the values of $\Delta\lambda_{20}$ and DSR_{max} are under study. It is expected that the value of $\Delta\lambda_{20}$ for the LR-2 application will be less than 1 nm.

† This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.

‡ For a 2-dB dispersion penalty.

The lower attenuation limit of 10 dB may require use of optical attenuators, lower maximum output power, increased minimum overload, or a combination thereof. Consequently, multisupplier compatibility is not guaranteed.

Table 4-7. IR OC-1 Optical Parameters

Parameter	IR-1	IR-2		Unit
Transmitter	MLM	MLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1260-1360	1430-1580	1430-1580	nm
$\Delta\lambda_{rms}$	23	4	NA	nm
$\Delta\lambda_{20}$	NA	NA	1	nm
SSR_{min}	NA	NA	30	dB
P_{Tmax}	-8	-8	-8	dBm
P_{Tmin}	-15	-15	-15	dBm
r_{emin}	8.2	8.2	8.2	dB
Optical Path				
System ORL _{min} *	NA	NA	NA	dB
DSR_{max}	96	317	NA	ps/nm
Attenuation	0-12	0-12	0-12	dB
Max. Reflectance between S and R	NA	NA	NA	dB
Receiver				
P_{Rmax}	-8	-8		dBm
P_{Rmin}	-28	-28		dBm
P_O	1	1		dB
Max. Receiver Reflectance	NA	NA		dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.

Table 4-8. IR OC-3 Optical Parameters

Parameter	IR-1	IR-2		Unit
Transmitter	MLM	MLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1261-1360	1430-1576	1430-1580	nm
$\Delta\lambda_{rms}$	7.7	2.5	NA	nm
$\Delta\lambda_{20}$	NA	NA	1	nm
SSR_{min}	NA	NA	30	dB
P_{Tmax}	-8	-8	-8	dBm
P_{Tmin}	-15	-15	-15	dBm
τ_{emin}	8.2	8.2	8.2	dB
Optical Path				
System ORL _{min} *	NA	NA	NA	dB
DSR_{max}	96	296	NA	ps/nm
Attenuation	0-12	0-12	0-12	dB
Max. Reflectance between S and R	NA	NA	NA	dB
Receiver				
P_{Rmax}	-8	-8		dBm
P_{Rmin}	-28	-28		dBm
P_O	1	1		dB
Max. Receiver Reflectance	NA	NA		dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.

Table 4-9. IR OC-12 Optical Parameters

Parameter	IR- 1	IR-2	Unit
Transmitter	MLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1293-1334†	1430-1580	nm
$\Delta\lambda_{rms}$	4.0 (2.5)‡	NA	nm
$\Delta\lambda_{20}$	NA	1	nm
SSR_{min}	NA	30	dB
P_{Tmax}	- 8	- 8	dBm
P_{Tmin}	- 15	- 15	dBm
r_{emin}	8.2	8.2	dB
Optical Path			
System ORL _{min} *	NA	24	dB
D_{SRmax}	46 (74)‡	NA	ps/nm
Attenuation	0-12	0-12	dB
Max. Reflectance between S and R	NA	- 27	dB
Receiver			
P_{Rmax}	- 8	- 8	dBm
P_{Rmin}	- 28	- 28	dBm
P_O	1	1	dB
Max. Receiver Reflectance	NA	- 27†	dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.
- † This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.
- ‡ Transmitters meeting the narrower spectral width objective are allowed a wider central wavelength range ($\lambda_{Tmin}-\lambda_{Tmax}$) of 1274-1356 nm.

Table 4-10. IR OC-48 Optical Parameters

Parameter	IR-1	IR-2	Unit
Transmitter	SLM	SLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1260-1360	1430-1580	nm
$\Delta\lambda_{rms}$	NA	NA	nm
$\Delta\lambda_{20}$	1	1**	nm
SSR_{min}	30	30	dB
P_{Tmax}	0	0	dBm
P_{Tmin}	-5	-5	dBm
r_{emin}	8.2	8.2	dB
Optical Path			
System ORL _{min} *	24	24	dB
D_{SRmax}	NA	**	ps/nm
Attenuation	0-12	0-12	dB
Max. Reflectance between S and R	-27	-27	dB
Receiver			
P_{Rmax}	0	0	dBm
P_{Rmin}	-18	-18	dBm
P_O	1	1	dB
Max. Receiver Reflectance	-27†	-27†	dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.
- ** Currently, there are no reliable methods for estimating chirp power penalties for SLM lasers; therefore, the values of $\Delta\lambda_{20}$ and D_{SRmax} are under study.
- † This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.

Table 4-11. SR OC-N Optical Parameters

Parameter	OC-1	OC-3	OC-12	OC-48	Unit
Transmitter Type	MLM/LED	MLM/LED	MLM/LED	MLM	
$\lambda_{Tmin}-\lambda_{Tmax}$	1260-1360	1260-1360	1261-1360	1266-1360	nm
$\Delta\lambda_{rms}$	80	40/80	14.5/35	4	nm
$\Delta\lambda_{20}$	NA	NA	NA	NA	nm
SSR_{min}	NA	NA	NA	NA	dB
P_{Tmax}	-14	-8	-8	-3	dBm
P_{Tmin}	-23	-15	-15	-10	dBm
r_{emin}	8.2	8.2	8.2	8.2	dB
Optical Path					
System ORL _{min} *	NA	NA	NA	24	dB
DSR_{max}	NA	18/25	13/14	12	ps/nm
Attenuation	0-7	0-7	0-7	0-7	dB
Max. Reflectance between S and R	NA	NA	NA	-27	dB
Receiver					
P_{Rmax}	-14	-8	-8	-3	dBm
P_{Rmin}	-31	-23	-23	-18	dBm
P_O	1	1	1	1	dB
Max. Receiver Reflectance	NA	NA	NA	-27†	dB

Notes:

- * For all applications, it is an objective that the optical system suffer a power penalty less than 1 dB in the presence of -8.5 dB reflectance.
- † This value is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations. For systems employing few or higher-performance optical components (e.g., a system with only two connectors), a -14 dB receiver reflectance may be considered acceptable.

4.3 Engineering of a Single-Mode Fiber Optic Transmission System

This section describes a methodology that may be used for engineering a single-mode fiber optic transmission system. The design approach this section describes is derived from EIA/TIA-559, *Single-Mode Fiber Optic System Transmission Design*, which is intended to apply to single-supplier systems operating up to about 0.5 Gb/s in the 1310-nm region over dispersion unshifted single-mode fiber. Extensions of the method to account for multi-supplier span designs and systems operation above about OC-12 are under study in TIA.

Parameters that allow the design of regenerator sections are detailed below. Some of the parameters in this section are defined relative to measurements specified in EIA/TIA documents (denoted by FOTP or OFSTP, some of which may be in draft form). Section 4.3.1 describes the transmission design information for terminal and regenerator station equipment, while Section 4.3.2 describes that for outside-plant cable. Worksheets 1 to 4 in Appendix B are examples of forms that could be used in gathering the transmission design information. Section 4.3.3 discusses how this information is used in the design and analysis of regenerator sections, which users or a system integrator can perform.

In this section, the transmission design parameters are specified by worst-case values.

- R4-14** [123] Suppliers shall provide worst-case values of transmission design parameters² requested as part of system documentation.
- CR4-15** [124] A BCC may require that suppliers guarantee the worst-case values over the lifetime of their system components.

4.3.1 Terminal Equipment Transmission Design Information

This section describes the transmission information for terminal and regenerator station fiber optic equipment that a supplier and system integrator must provide. Worksheets 1 through 3 contain a summary of terminal equipment transmission design information.

- R4-16** [125] The supplier shall provide general terminal equipment information in Worksheet 1.
- R4-17** [126] The supplier and the system integrator shall provide terminal equipment parameters under normal operating and short-term emergency conditions in Worksheets 2 and 3, respectively.
- R4-18** [127] If the terminal equipment has many options, the information of the worksheets shall be provided for each option.

2. Some of the parameters are to be provided by the "system integrator" responsible for the overall system design. This may or may not be the supplier.

4.3.1.1 General System Information

Terminal Equipment Identification – The terminal equipment identification uniquely characterizes the type of product and provides a traceable indicator for determining the product specifications, features, issue or revision, and manufacturer (e.g., CLEI™).

Optical Line Rate – $N \times 51.840$ Mb/s, where $N = 1, 3, 12, 24, 48$ or 192 .

4.3.1.2 Transmitter Information

The unit providing the transmitter function is identified by a unique descriptor from which the following information can be determined, using the appropriate documentation:

- Manufacturer
- Terminal equipment association
- System design application (e.g., single-mode long-reach)
- Operating wavelength
- Output power level
- Source type
- Optical device temperature controller
- FDA classification (e.g., Class I or Class II)
- Manufacturer product change designation (e.g., issue or revision).

Generic transmitter requirements are in TA-NWT-001385, *Generic Requirements for Optoelectronic Devices in Fiber Optic Systems*.

Optical Source Type – The optical source type is characterized by identifying, as a minimum:

- Type of device (e.g., LED, edge emitting LED, SLM laser, or MLM laser)
- Material composition of source (e.g., InGaAs)
- Generic device structure (e.g., Distributed Feedback [DFB]).

Transmitter Connector – The transmitter connector is the optical connector provided at the output of the transmitter that attaches to the transmitter pigtail. The transmitter connector description, as a minimum, includes:

- Connector manufacturer
- Connector type (e.g., Biconic, FC)
- Connector model number
- Connector classification (multimode, single-mode)
- Mating connector model number.

Generic connector requirements are in GR-326-CORE.

Transmitter Pigtail – The identification of the transmitter pigtail includes the following information (see EIA/TIA-492, *Generic Specification for Optical Waveguide Fiber*):

- General fiber type
- Class of fiber
- Mode field diameter.

Maximum Optical Reflection (OR_{max}) – The total reflected optical power (in dB) that a transmitter can accommodate and maintain its stated performance (at 10^{-10} BER). Section 4.2 contains the lower bounds to OR_{max} .

The system integrator specifies the nominal central wavelength $\Delta\lambda_{Tnom}$, transmitter central wavelength range ($\Delta\lambda_{Tmin}$, $\Delta\lambda_{Tmax}$) and transmitter power (P_T), as Section 4.2.5 defines.

4.3.1.3 Receiver Information

4.3.1.3.1 General Information

The unit providing the receiver function is identified by a unique descriptor from which the following information can be determined, using the appropriate documentation:

- Manufacturer
- Terminal equipment association
- System design application (e.g., single-mode long-reach)
- Receiver performance specifications
- Detector type
- Optical device temperature controller
- Manufacturer product change designation (e.g., issue or revision).

Generic receiver requirements are in TA-NWT-001385.

Optical Detector Type – The optical detector type is characterized by identifying as a minimum:

- Device type (e.g., Positive-Intrinsic-Negative [PIN], Avalanche Photodiode [APD])
- Material composition of detector (e.g., Ge, Si).

Receiver Connector – The receiver connector is the optical connector provided at the input to the receiver that attaches to the receiver pigtail. The receiver connector description, as a minimum, includes:

- Connector manufacturer
- Connector type (e.g., Biconic, FC)
- Connector model number

Connector classification (multimode, single-mode)
Mating connector model number.

Generic connector requirements are in GR-326-CORE.

Receiver Pigtail – The identification of the receiver pigtail includes the following information (see EIA/TIA-492):

General fiber type
Class of fiber
Mode field diameter.

4.3.1.3.2 Transmission Properties

The transmission parameters to be specified for the receiver unit are:

Receiver Sensitivity (P_R) – The worst-case value of the input optical power (dBm) to the receiver (on the line side of the receiver module connector), specified as P_{R1} for standard operating conditions or P_{R2} for extended operating conditions, that is necessary to achieve the manufacturer-specified BER as measured using the procedure in OFSTP-3, *Fiber Optic Terminal Receiver Sensitivity and Maximum Receiver Input Power*.

The receiver sensitivity value specified includes the following performance degradation factors combined in a worst-case fashion:

- Manufacturing variations with temperature and aging drifts, including any degradation of the receiver connector
- Maximum transmitter power penalty resulting from the use of a transmitter with a worst-case extinction ratio r_e when operated under standard operating conditions (P_{R1}) or extended operating conditions (P_{R2})
- Maximum transmitter power penalty resulting from the use of a transmitter with a worst-case rise/fall time when operated under standard operating conditions (P_{R1}) or extended operations (P_{R2})
- Maximum transmitter power penalty resulting from worst-case optical return loss at point S.

The receiver sensitivity does not include power penalties associated with dispersion (pulse broadening), reflections, or jitter. OFSTP-11, *Measurement of Single Reflection Power Penalty for Fiber Optic Terminal Equipment*, contains a procedure to determine the power penalty associated with worst-case reflections at the line side of the transmitter connector (Point S).

Dispersion Power Penalty (P_D) – The maximum power penalty (dB) associated with the worst-case increase in receiver input optical power level to account for the total pulse distortion due to ISI, Mode Partition Noise (MPN), and laser chirp at the specified bit rate,

BER of 10^{-10} , and Maximum Dispersion (D_{SRmax}), when operated under standard operating conditions (P_{D1}) or extended operating conditions (P_{D2}).

OFSTP-10, *Measurement of Dispersion Power Penalty in Single-Mode Systems*, contains a procedure for measuring dispersion power penalty.

P_D equals or is less than the value of P_O for each application that Tables 4-3 through 4-11 describe.

Maximum Dispersion (D_{SRmax}) – The maximum dispersion (ps/nm), due to fiber length between points S and R (see Figure 4-1) that can be accommodated by a transmitter-receiver pair to meet the 10^{-10} BER performance objective, when operated under standard operating conditions (D_{SRmax1}) or extended operating conditions (D_{SRmax2}).

D_{SRmax} may exceed the values that Tables 4-3 through 4-11 specify.

Receiver Overload Point (R_{max}) – The maximum value of the input optical power (dBm) to the receiver (on the line side of the receiver module connector, point R), when operated under standard operating conditions R_{max1} , or extended operating conditions R_{max2} , that the receiver will accept and maintain a 10^{-10} BER as measured using the procedure in OFSTP-3. This value does not include the effects of a removable optical attenuator that may be needed to meet the maximum average received power values P_{Rmin} , specified for the applications in Tables 4-3 through 4-11, that could be placed on the receiver side of point R (Figure 4-1) and, consequently, be regarded as part of the receiver.

R_{max} may be greater than the value P_{Rmax} .

The receiver parameters P_R , P_D , D_{SRmax} , and R_{max} are provided for BER of 10^{-10} as Worksheets 2 and 3 show. Optionally, the BCC may request these parameters for other BER values.

4.3.1.4 Attenuators

The supplier provides a full description of the attenuators that are to be used with the system, if needed. The specifications, as a minimum, include:

U_{att} = Insertion loss (dB).

OR_{att} = Worst-case value of attenuator reflectance (dB).

Attenuator criteria are in GR-910-CORE.

4.3.1.5 Wavelength Division Multiplex (WDM) Device

If a WDM device is offered, the supplier specifies the manufacturer, model number, number of channels, and loss:

U_{WDM} = Worst-case value of the all-inclusive loss (dB) associated with WDM equipment (at both ends), including all insertion and additional connector losses as well as other degradations. The allocations must include the effects of temperature, humidity, and aging.

The loss corresponds to the transmitter wavelength stated in Worksheets 2 and 3. WDM criteria are in GR-1209-CORE.

4.3.1.6 Safety Margin

M = Safety margin, in dB, for unexpected losses, to be determined by the system integrator for a specific application.

M does not include penalties for expected losses and degradations (e.g., laser aging or reflections) because these effects are already included in the appropriate transmission parameters, according to the definitions in this GR. The specifications in Section 4.2 assume $M = 0$. When a BCC system integrator desires, a non-zero value of M may be used. In this case, the value is to be included in determining the effective attenuation as Tables 4-3 through 4-11 specify. Selecting a non-zero value of M for the system design with transmitters meeting the specifications in Tables 4-3 through 4-11 and receivers meeting the minimum objective value for receiver overload P_{Rmax} in the tables, may require optical attenuators to protect the receiver from damage caused by optical power overload.

4.3.1.7 Connectors

The supplier specifies the following connector information:

- Connector type (e.g., Biconic, FC)
- Manufacturer
- Model number
- Connector classification (multimode, single-mode).

Also, the supplier specifies these connector parameters:

U_{con} = Worst-case value of connector loss (dB).
 OR_{con} = Worst-case value of connector reflectance (dB).

Generic connector requirements are in GR-326-CORE.

Connector Variation – Connector variation is the maximum value (dB) of the difference in insertion loss between mating optical connectors of the same type and model, from the same manufacturer.

The system integrator specifies the following:

N_{con} = Number of single-mode to single-mode connectors. This is the number recommended by a system integrator for a typical point-to-point regenerator section. This should not include the transmitter unit or receiver unit connectors, because they are already accounted for in P_T and P_R , respectively.

4.3.1.8 Station Cable

Station cable represents the optical fiber cable that is used within a building environment to connect the outside plant optical fiber cable to the optical fiber system terminal equipment. The station cable may provide this optical path by means of some form of optical patch panel that allows optical path rearrangement to the outside plant fibers.

The supplier provides the following information:

- Manufacturer
- General fiber type (see EIA/TIA-492)
- Class of fiber (see EIA/TIA-492).

Interconnection-related parameters:

- Nominal mode field diameter and tolerance
- Nominal cladding diameter and tolerance
- Maximum cladding ovality
- Maximum core/cladding concentricity error.

The interconnection-related parameters are needed to calculate the connection losses of field-installed splicers and connectors.

Also, the supplier specifies the following transmission parameters:

U_{SM} = Worst-case end-of-life loss (dB/km) of single-mode regenerator station cable.

λ_{cc} = Cable cutoff wavelength (refer to EIA/TIA-455-170). The cutoff wavelength of the fiber jumper cable must be below the minimum value of the transmitter central wavelength.

The system integrator specifies the following:

I_{SM} = Total length in km, on both ends of a regenerator section, of single-mode regenerator station cable.

This document does not include the specification of optical and physical parameters for fiber. GR-20-CORE and GR-409-CORE, *Generic Requirements for Premises Fiber Optic Cable*, contain such criteria.

4.3.2 Cable Transmission Design Information

For a given cable type, the cable supplier provides two categories of cable transmission information:

1. Parameters for specific applications. These are specified at the time of initial installation.
2. Global loss and chromatic dispersion characteristics. Because these may not be guaranteed values, the user should use them for initial feasibility studies only, and should measure the parameters for a specific upgrade. If a known upgrade is planned at the time of initial installation, the required parameters should be specified.

The following sections explain each category.

4.3.2.1 Parameters for a Specific Application

This section discusses the cable transmission parameters to be provided by the system integrator and by the suppliers. They are summarized in Worksheet 4. The parameters discussed in this section are to be given as worst-case values.

The system integrator specifies the following:

- Type of application (e.g., aerial, buried, or underground)
- Temperature range (e.g., -7°C to 40°C for underground and -40°C to 77°C for aerial environments)
- The cabled fiber reel length l_R (in km)
- The nominal central wavelength λ_{Tnom} and central wavelength range (λ_{Tmin} to λ_{Tnom}) corresponding to the terminal equipment to be used
- The type of splice, if appropriate
- Splice loss information:

U_S = Maximum allowable splice loss (in dB/splice) at 23°C .

U_{ST} = Effect of temperature on splice loss (in dB/splice) at the worst-case temperature conditions, over the specified cable operating temperature range. If U_S already includes corrections for U_{ST} , then U_{ST} will be zero.

The supplier specifies the following:

- Designation of the cable
- Maximum cable cutoff wavelength, λ_{cc} , (in nm) as per EIA/TIA-455-170, with cable deployment conditions as Figure 4-4 shows.

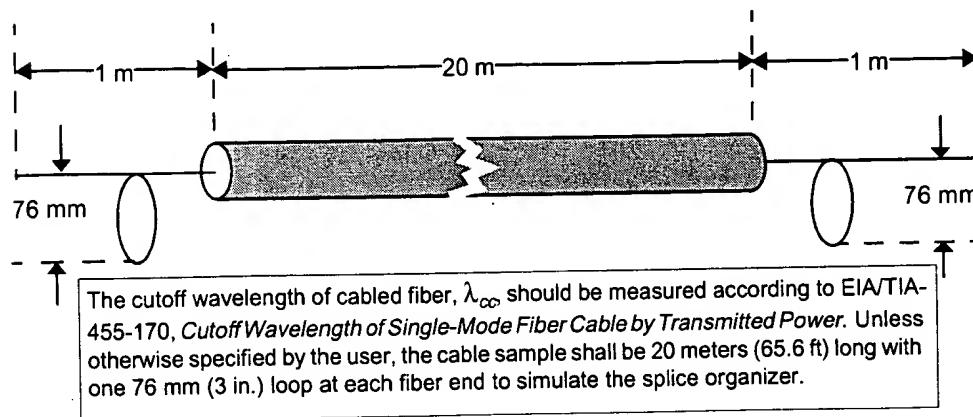


Figure 4-4. Cable Configuration for Cabled Fiber Cutoff Wavelength Measurement

The cutoff wavelength of a cabled optical fiber demarcates the wavelength region above which the fiber supports propagation of only a single mode and below which multiple modes are supported. Operation below the cutoff wavelength may result in modal noise, modal distortion (increased pulse broadening), and improper operation of connectors, splices, and WDM couplers. For these reasons, the system operating wavelength range, dictated by the transmitter central wavelength range, described in Section 4.2.4.1, must be greater than the maximum allowed cutoff wavelength to ensure the system is operating entirely in the fiber's single-mode regime. In general, the highest value of cabled fiber cutoff wavelength, λ_{ccmax} , will be found in the shortest installation or repair cable length. A criterion that ensures a system is free from high cutoff wavelength problems is:

$$\lambda_{ccmax} < \lambda_{Tmin} \quad (4-8)$$

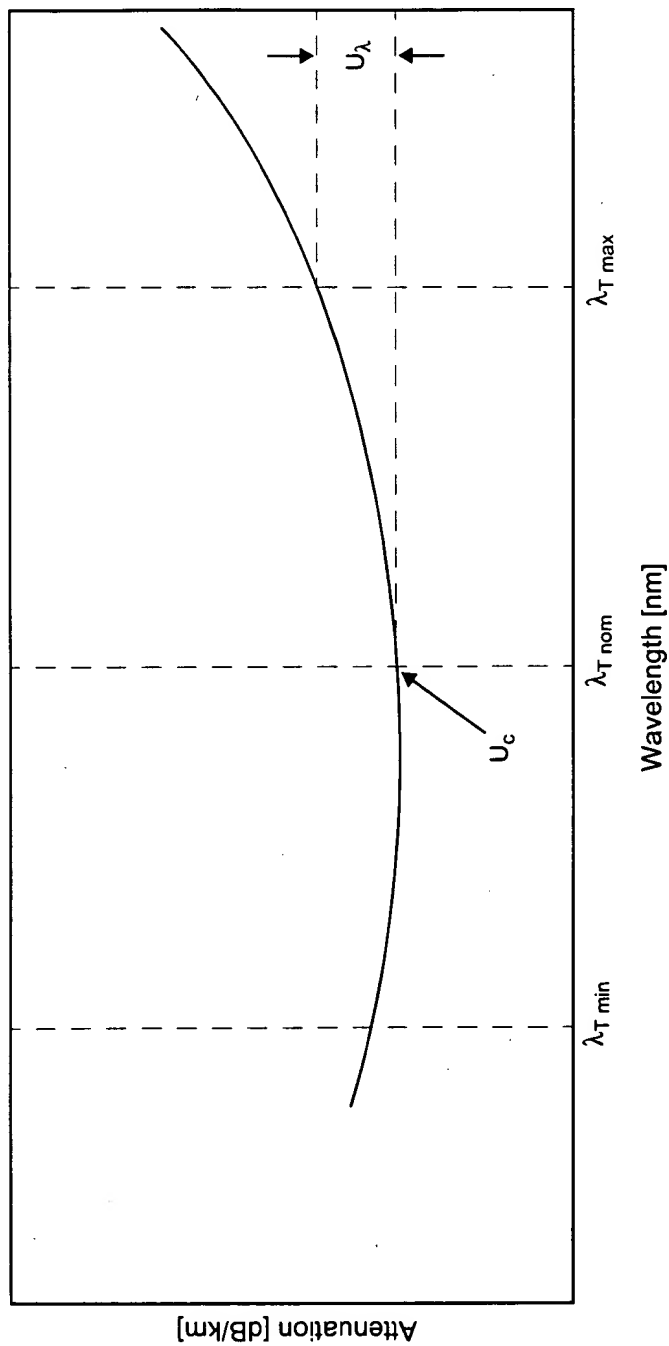
where Section 4.2.4.1 defines λ_{Tmin} .

- Cable loss parameters:

U_c = Worst-case end-of-life cable loss (in dB/km at 23°C) at the transmitter's nominal central wavelength λ_{Tnom} . This includes splicing loss caused by the fiber or cable manufacturing process.

U_{λ} = The largest increase in cable loss (in dB/km at 23°C) above U_c which occurs over the transmitter's central wavelength range (λ_{Tmin} to λ_{Tmax}). Figure 4-5 illustrates the determination of this.

U_{cT} = Effect of temperature on end-of-life cable loss (in dB/km) at the worst-case temperature conditions over the specified cable operating temperature range.



Note: This figure is only an example. In actuality, U_λ may occur at $\lambda_{T \max}$, $\lambda_{T \min}$, or another wavelength between $\lambda_{T \min}$ and $\lambda_{T \max}$ relative to $\lambda_{T \text{ nom}}$.

Figure 4-5. Wavelength Dependent Attenuation Characteristics

- Dispersion parameters:

$\lambda_{0 \min}$, $\lambda_{0 \max}$ = Minimum and maximum values of zero-dispersion wavelengths.

$S_{0 \max}$ = Maximum value of the zero-dispersion slope [the dispersion slope in ps/(nm² × km) at the zero-dispersion wavelength].

D_{\max} = Absolute value of the worst case of chromatic dispersion in ps/(nm × km) over the transmitter central wavelength range and over the allowed variation in the cable type's dispersion.

- Interconnection related parameters: The supplier specifies the nominal mode field diameter and tolerance, the nominal cladding diameter and tolerance, the maximum cladding ovality, and the maximum core/cladding concentricity error. These parameters are needed to calculate the connection losses of field-installed splices and connectors.

If a future system upgrade is planned, the user should specify the performance required for the upgrade. The supplier furnishes the information of Worksheet 4 for the original application and the future upgrade.

4.3.2.2 Global Fiber Parameters

A future application of an installed cable may arise that was not anticipated when the cable was purchased (e.g., an unforeseen upgrading of terminal equipment, rerouting of traffic, or restoration). In such cases, global fiber characteristics may be helpful for indicating to the user the feasibility of the cable's proposed application.

The fiber global loss and dispersion characteristics are curves that depict these parameters as a function of wavelength. They are to be provided over the fiber's useful wavelength regions, for example, from 1260 to 1360 nm and from 1430 to 1580 nm. The global characteristics should show typical values because the only worst-case values are those specified when the cable was purchased. Thus, if the global characteristics indicate the new application to be feasible, measurements still must be taken to verify that the cable will perform properly.

For the fiber global loss characteristics, the supplier provides a graph similar to Figure 4-6, which shows the loss at 23°C and the maximum temperature effect over the cable's operating temperature range, U_{CT} , as a function of wavelength. Also, the system integrator describes the maximum allowable splice loss, U_S , at 23°C and the maximum temperature effect on splices, U_{ST} , as a function of wavelength. These characteristics are provided for underground, buried, and aerial applications, over the temperature ranges as the user specifies.

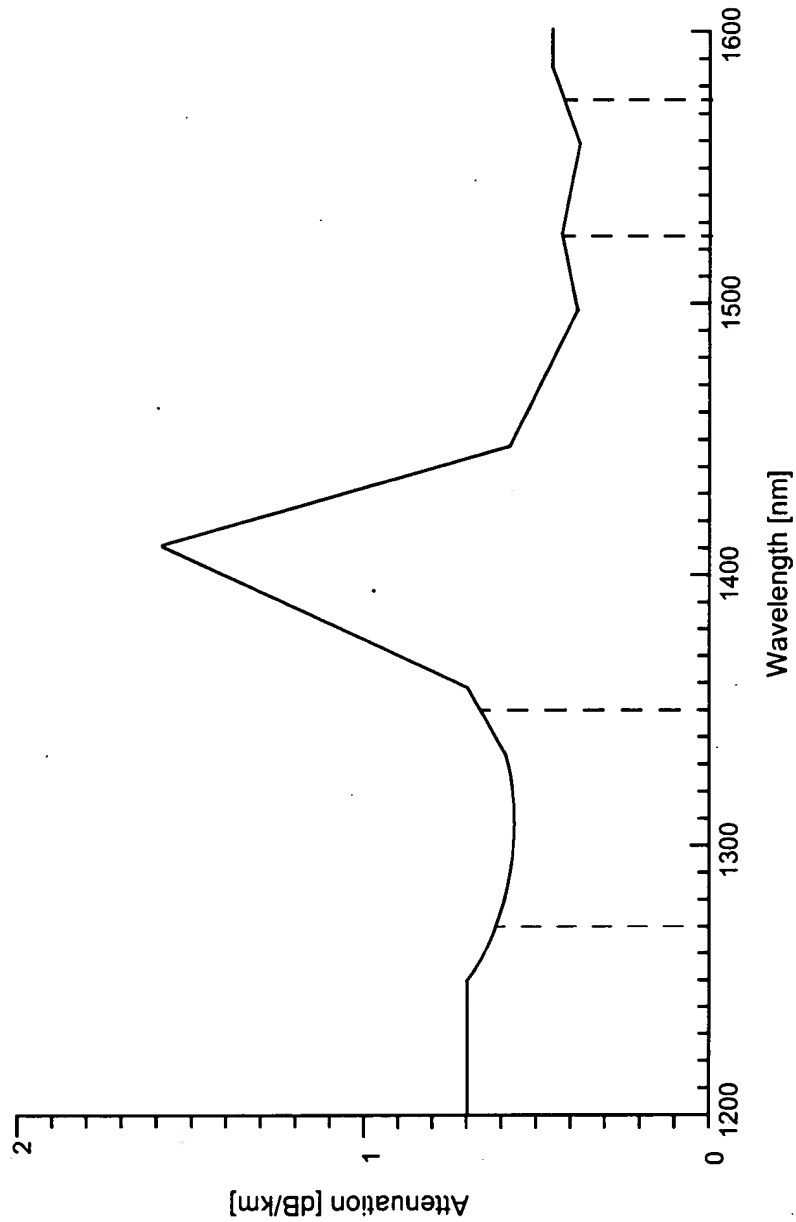


Figure 4-6. Global Fiber Attenuation Characteristics

For the global dispersion characteristic, the data that Section 4.3.2.1 and Worksheet 4 reference can be used with the following formulas to approximate the fiber's chromatic dispersion coefficient D , in ps/(nm × km), as a function of wavelength. The extreme limits, $D_1(\lambda)$ and $D_2(\lambda)$, are given as:

$$D_1(\lambda) = \frac{S_{o \max}}{4} \left(\lambda - \frac{\lambda_{o \min}^4}{\lambda^3} \right) \quad (4-9)$$

$$D_2(\lambda) = \frac{S_{o \max}}{4} \left(\lambda - \frac{\lambda_{o \max}^4}{\lambda^3} \right) \quad (4-10)$$

Figure 4-7 shows the respective curves.

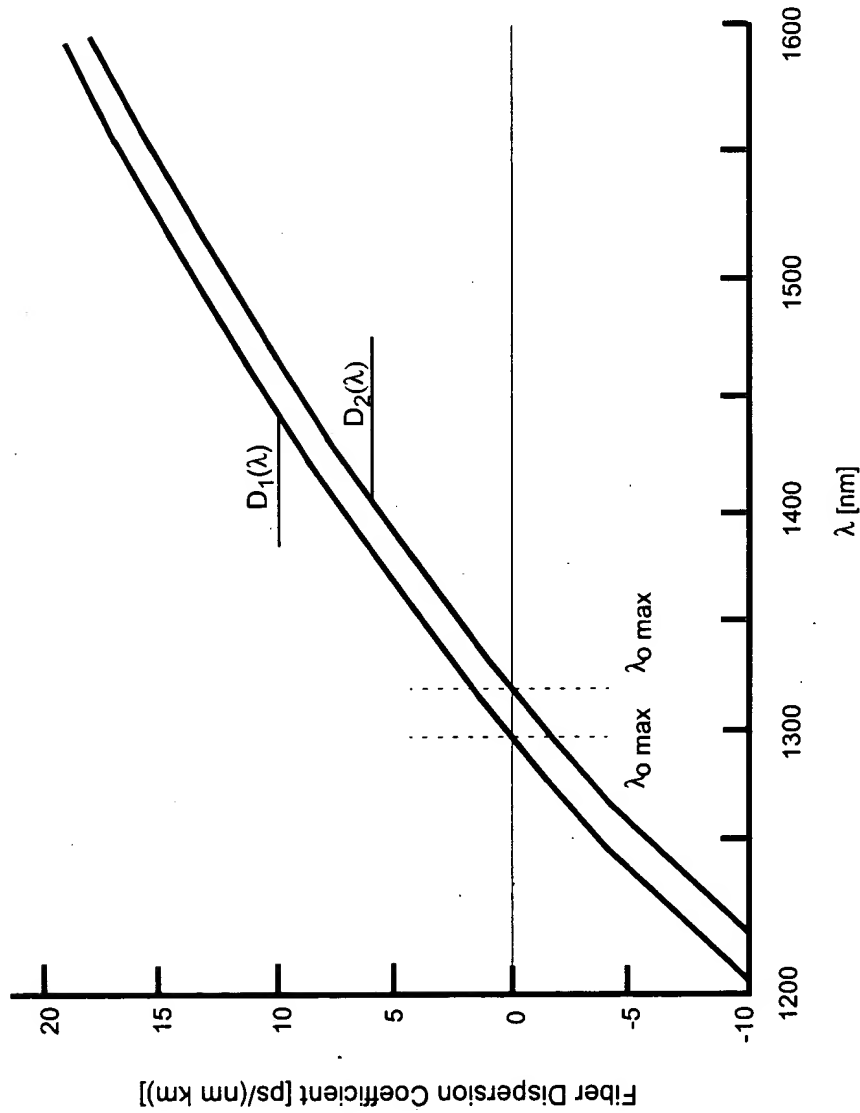


Figure 4-7. Global Fiber Dispersion Characteristics

4.3.3 Fiber Optic System Transmission Design and Analysis

4.3.3.1 Design Approach

This section describes a methodology for engineering the route layout for a single-mode fiber optic transmission system. The methods address the design of a regenerator section and are intended to ensure that all installed fiber paths that meet cable specifications in every regenerator section will be usable at certain proposed transmission rates. A regenerator section, defined in Figure 4-8, is composed of: 1) regenerators, connectors, station cables, and a fiber interconnection feature at each regenerator station; and 2) fiber cable sections, joined by splices, between regenerator stations. The fiber interconnection feature is the interface between the regenerator station equipment and the fiber cable. It provides connector access to each fiber in the cable for maintenance and restoration purposes.

This section describes how the transmission parameters of Sections 4.3.1 and 4.3.2 are used in the design and analysis of regenerator sections. The computations are based on two constraint relations, one on the system's loss budget and the other on dispersion tolerance. Each of these is discussed. Finally, the design and analysis methodology is discussed.

Selecting a non-zero value of M for the system design with transmitters meeting the specifications in Tables 4-3 through 4-11 and receivers meeting the minimum objective value for receiver overload P_{Rmax} in the tables, may require optical attenuators to protect the receiver from damage caused by optical power overload.

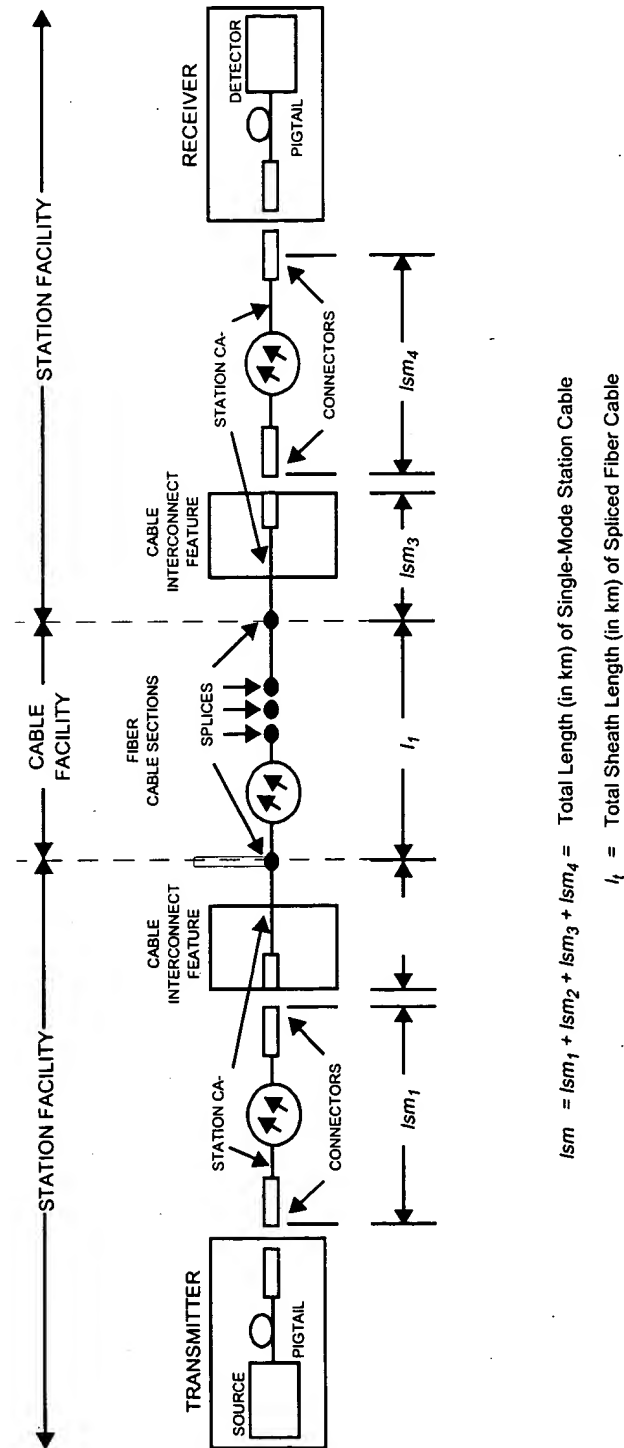


Figure 4-8. Engineering of a Fiber Optic Regenerator Section

4.3.3.2 Loss Budget Constraint

Figure 4-9 shows a model for transmission in a regenerator section including transmitter power, receiver sensitivity, and various sources of loss. For route layout design, it is convenient to divide the transmission path at the fiber distributing frame into terminal or regenerator station equipment and spliced fiber cable. The system gain, G , of terminal or regenerator station equipment can then be determined, and the loss, L , introduced by the fiber path between regenerator interfaces, can be determined. If L is less than or equal to G , then the embodied transmission objectives can be attained. An attenuator may be needed if L is much smaller than G .

In the following equations, the parameters that enter into G and L are specified by their worst-case values.

The equation for the system gain of the terminal or regenerator station equipment is

$$G = P_T - P_R - P_D - M - U_{WDM} - l_{SM} U_{SM} - N_{con} U_{con} \quad (4-11)$$

where the terms are as Section 4.3 defines.

The equation for the loss of the fiber cable within a regenerator section is

$$L = l_t(U_c + U_{cT} + U_\lambda) + N_S(U_S + U_{ST}) \quad (4-12)$$

where:

l_t = Total sheath length of spliced fiber cable in km. This value may include an allowance for cable repair purposes.

N_S = Number of splices in fiber cable of length l_t (including those at the fiber distributing frame at both ends of the regenerator section), plus any allowance for additional splices for cable repair purposes.

The remaining terms are as defined in Section 4.3.2.1.

Section 4.3.3.4 discusses transmission analysis and design calculations using G and L .



Figure 4-9. Regenerator Section, Sources of Loss

4.3.3.3 Dispersion Limited Length

In general, for single-mode systems operating at rates less than approximately 0.5 Gb/s, regenerator section length may be expected to be limited by loss and not by dispersion. At higher bit rates, however, the length may be limited by dispersion, which arises from the combined effects of chromatic dispersion, mode-partition noise, and laser chirp. Thus, it is desirable to check whether a regenerator section's length may be limited by dispersion.

A regenerator section will not be limited by dispersion as long as the following holds:

$$D(\lambda_t) \times l \leq D_{SRmax}, \quad (4-13)$$

where:

- $D(\lambda_t)$ = fiber chromatic dispersion coefficient evaluated at λ_t in ps/(nm × km)
- l = fiber length in km
- D_{SRmax} = maximum dispersion (in ps/nm) as Section 4.3.1.3.2 defines. For dispersion-limited systems, Tables 4-3 through 4-11 specify maximum values for D_{SRmax} .

Taking into account the variation in the fiber chromatic dispersion coefficient over the transmitter's central wavelength range, the dispersion limited length is given by

$$l_D = \frac{D_{SRmax}}{D_{max}}, \quad (4-14)$$

where:

- D_{max} = absolute value of the worst case of chromatic dispersion coefficient in ps/(nm × km) over the specified range of the transmitter central wavelength and over the specified variation in the cable type's dispersion.
- l_D = dispersion-limited length.

The effect of dispersion is accounted for in the equation for G (Equation 4-11, page 4-39) via the dispersion power penalty P_D in dB.

4.3.3.4 Design and Analysis Methodology

The transmission design and analysis methodology for fiber optic systems focuses on the individual sections making up the overall system. The design and analysis calculations are performed using specific combinations of terminal equipment and cable parameters.

The design algorithm deals with a proposed fiber optic system route with given distances between terminal and regenerator locations. This document assumes SONET regenerator sections to be designed for 10^{-10} BER. However, individual BCCs may have a different BER requirement. For each regenerator section, one must determine whether a particular combination of terminal equipment and cable can work together (at the desired BER) without exceeding any loss or dispersion limitation. The steps involved in the design algorithm are:

- a. From Worksheet 2 or 3, select the receiver parameters P_R , P_D , and D_{SRmax} corresponding to the desired BER.
- b. Calculate G and L (Equations 4-11 and 4-12, page 4-39). If $G - L \geq 0$, then the regenerator section is not limited by loss.
- c. Calculate the end-to-end regenerator section dispersion $D_{max} \times l$. If this is less than or equal to D_{SRmax} , then the regenerator section is not limited by dispersion (Equation 4-13, page 4-41).

If the calculations from steps b and c are satisfactory, the terminal equipment-cable combination is acceptable.

The analysis algorithm deals with calculating maximum regenerator section lengths. Such calculations may be used in comparing various combinations of terminal equipment and cable. It is suggested that the maximum regenerator section lengths be calculated as the distance between splices is varied, ranging from 0.5 km up to the length of a reel of cable. The steps involved in the analysis algorithm are:

- Select a set of receiver parameters P_R , P_D , and D_{SRmax} corresponding to the desired BER value from Worksheet 2 or 3.
- Calculate the length at which $G = L$ (Equations 4-11 and 4-12, page 4-39).
- Calculate the dispersion-limited length (Equation 4-14, page 4-41).
- Take the maximum regenerator section length to be the smaller of the two lengths calculated in steps b and c above.

4.4 Electrical Interface Specifications

In some applications, it may be desirable to interconnect SONET NEs using standard electrical interfaces. Central office engineering considerations limit the feasibility of electrical interfaces to the STS-1 and STS-3 levels of the SONET hierarchy. These interfaces are based on the following:

- One coaxial line is used for each direction of transmission. Reference cable is 75- Ω coaxial cable with tinned copper shield (AT&T Technologies, Inc. 728A or equivalent).
- Since the SONET signal is scrambled (Section 5.1.3), the electrical interface signal can be adequately modeled as one having a random occurrence of ones and zeros with mark probability nominally equal to 0.5. This allows specification of the power level at the interface in terms of a wideband power measurement rather than a narrow-band measurement power (as is used for DS1 interfaces). A wideband measurement is simpler than a narrow-band measurement because it does not require tightly controlled bandpass filters with specific roll-off characteristics.³
- The pulse shape at the interface is specified by an eye diagram mask, which is compatible with the pulse template that is also specified. Signals meeting the pulse template specification should also meet the eye diagram mask specification. In rare cases, however, a signal with a compliant eye diagram may not meet the pulse template specification. It is therefore recommended that the electrical interfaces be designed and tested against the pulse template specification. The eye diagram mask specification is included to allow a visual check of the waveform at the interface to identify obvious, overall waveform defects (including pulse imbalance) during trouble isolation testing. ANSI T1.102 describes procedures for checking conformance with the pulse and eye diagram masks.

CR4-19 [128v2] A SONET NE may be required to support STS-1 electrical interfaces.

CR4-20 [910] A SONET NE may be required to support STS-3 electrical interfaces.

R4-21 [129] If a SONET NE supports STS-1 electrical interfaces, the following apply:

- The transmitter shall generate an interface signal that meets the criteria in Table 4-12 for the entire range of interconnect cable lengths of 0 to 450 ft between the transmitter and the interface.

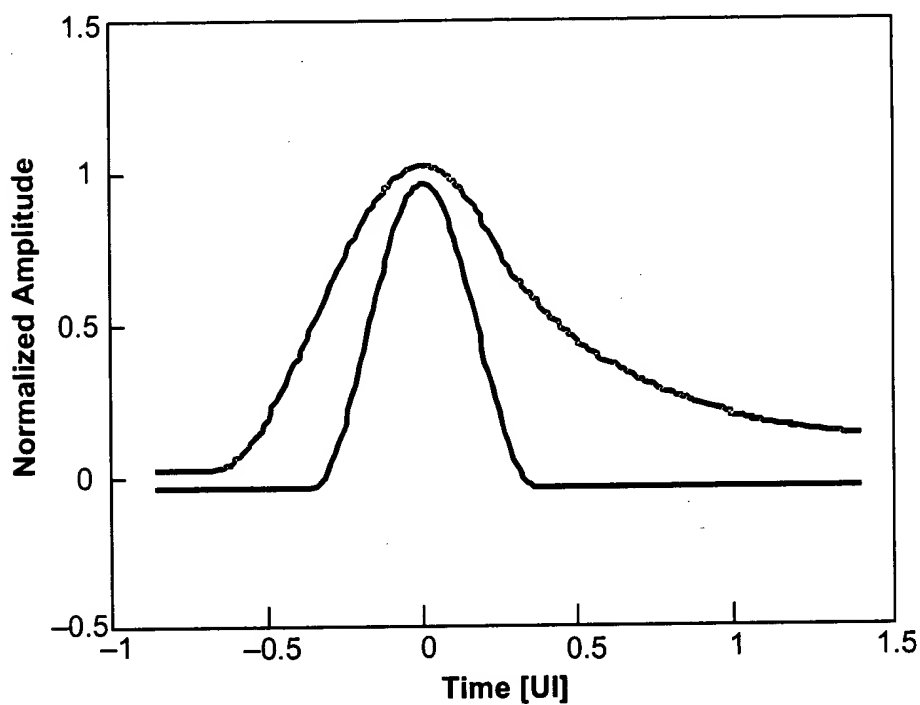
3. In the DS_n interface criteria in GR-499-CORE, requirements are given for both power levels and pulse amplitudes. Similarly, for STS-3 electrical signals, a requirement is given for the pulse amplitude at transmitter, and another requirement is given for the power level at the interface. For STS-1 electrical signals, only an interface power level requirement is given. A pulse amplitude requirement is not necessary for conformance testing purposes; however, in some situations (e.g., during trouble-shooting procedures) it may be useful (and simple) to measure the STS-1 pulse amplitude. An STS-1 electrical signal that meets the power level requirement is expected to have pulse amplitudes (at the interface) in the range of approximately 0.45 V to 0.9 V. Note that this range (0.45 to 0.9) is a conservative estimate calculated for signals having "worst-case" pulse shapes. Therefore, signals with pulse amplitudes somewhat outside this range may still meet the power level requirement.

- ...
- The receiver shall accept any interface signal that conforms to the criteria in Table 4-12 (at the interface), and that propagates through a jumper cable (if used, may be up to 27 ft) and an additional length of interconnect cable, which can also be up to 450 ft.

Note that a SONET NE's STS-1 electrical interfaces may switch in Line Build Out (LBO) circuits to maintain the equivalent interconnect distance between 225 and 450 ft where the actual cable length is less than 225 ft.

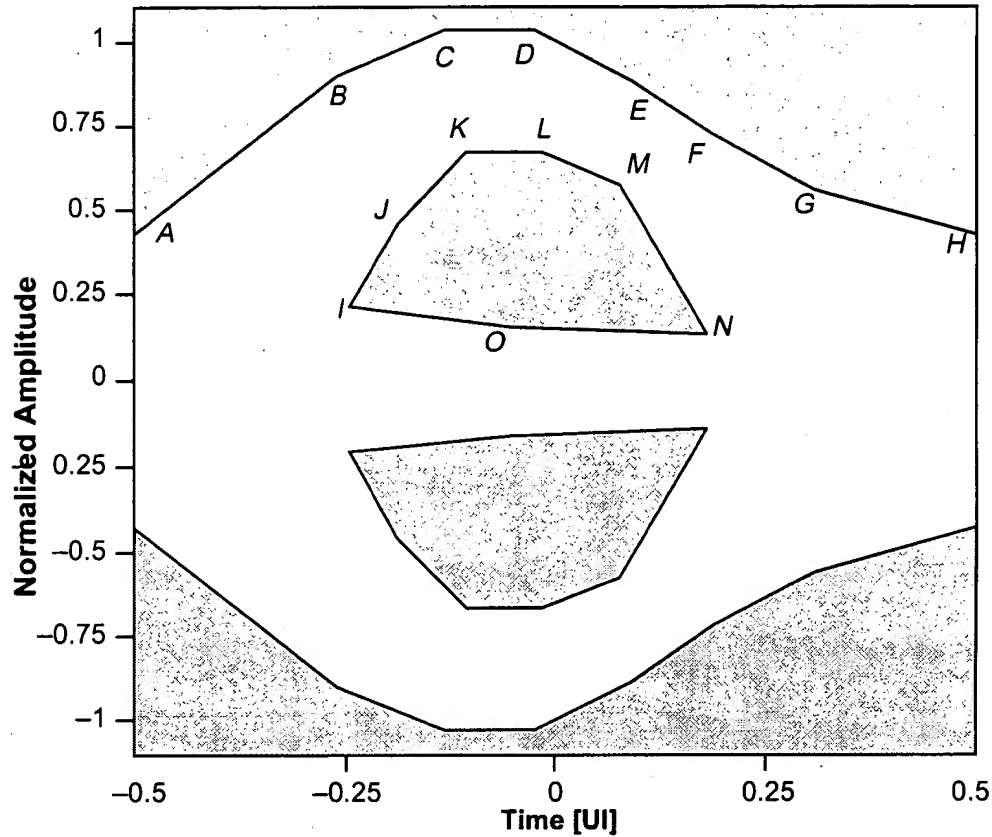
Table 4-12. STS-1 Electrical Interface Criteria

Line Rate	51.840 Mb/s (synchronized to the NE clock; see Section 5.4)
Line Code	Bipolar with Three-Zero Substitution (B3ZS)
Impedance	A resistive test load of 75Ω ($\pm 5\%$) is used at the interface for the evaluation of pulse shape and the electrical parameters specified below.
Pulse Shape	The shape of every pulse that approximates an isolated pulse (i.e., a pulse preceded by two zeros and followed by one or more zeros) shall fit within the limits of the pulse template in Figure 4-10. In addition, an eye diagram of the interface signal shall fit within the limits of the eye diagram mask in Figure 4-11.
Power Level	The wideband power level measured using a low-pass filter with a flat pass band and a 3-dB cutoff frequency of at least 207.36 MHz shall be between -2.7 dBm and +4.7 dBm.
DC Offset	There shall be no dc power flow across the interface.



	Time axis range [UI]	Normalized amplitude equation
Upper curve	$-0.85 \leq T \leq -0.68$	0.03
	$-0.68 \leq T \leq 0.26$	$0.5 \{1 + \sin [(\pi/2)(1 + T/0.34)]\} + 0.03$
	$0.26 \leq T \leq 1.40$	$0.1 + 0.61 e^{-2.4(T-0.26)}$
Lower curve	$-0.85 \leq T \leq -0.36$	-0.03
	$-0.36 \leq T \leq 0.36$	$0.5 \{1 + \sin [(\pi/2)(1 + T/0.18)]\} - 0.03$
	$0.36 \leq T \leq 1.40$	-0.03

Figure 4-10. STS-1 Electrical Interface Pulse Mask



Outer region corner points			Inner region corner points		
Point	Time	Amplitude	Point	Time	Amplitude
A	-0.5	0.426	I	-0.245	0.214
B	-0.261	0.904	J	-0.187	0.455
C	-0.136	1.03	K	-0.104	0.67
D	-0.028	1.03	L	-0.017	0.67
E	0.094	0.883	M	0.077	0.581
F	0.187	0.723	N	0.18	0.14
G	0.31	0.566	O	-0.054	0.16
H	0.5	0.426			

NOTE — Both inner and outer regions are symmetric about zero-amplitude axis

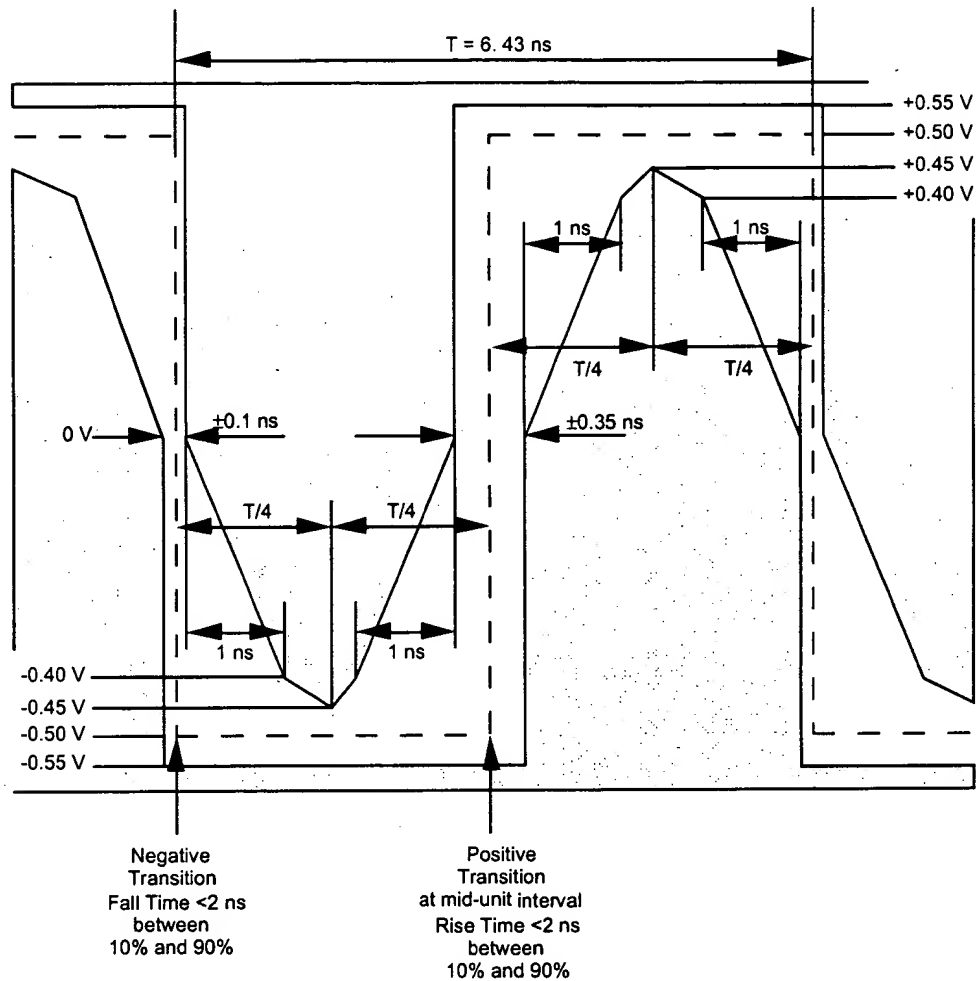
Figure 4-11. STS-1 Electrical Interface Eye Diagram Mask

R4-22 [130] If a SONET NE supports STS-3 electrical interfaces, the following apply:

- ... • The signal at the output of the transmitter (at the cable terminal) shall have a nominal rectangular pulse shape with a peak amplitude of 0.5 V (+10%) and maximum rise/fall times of 2 ns, as Figures 4-12 and 4-13 show.
- ... • The transmitter shall generate an interface signal that meets the criteria in Table 4-13 for the entire range of interconnect cable lengths of 0 to 225 ft between the transmitter and the interface.
- ... • The receiver shall accept any interface signal that conforms to the criteria in Table 4-13 (at the interface), and that propagates through a jumper cable (if used, may be up to 27 ft) and an additional length of interconnect cable, which can also be up to 225 ft.

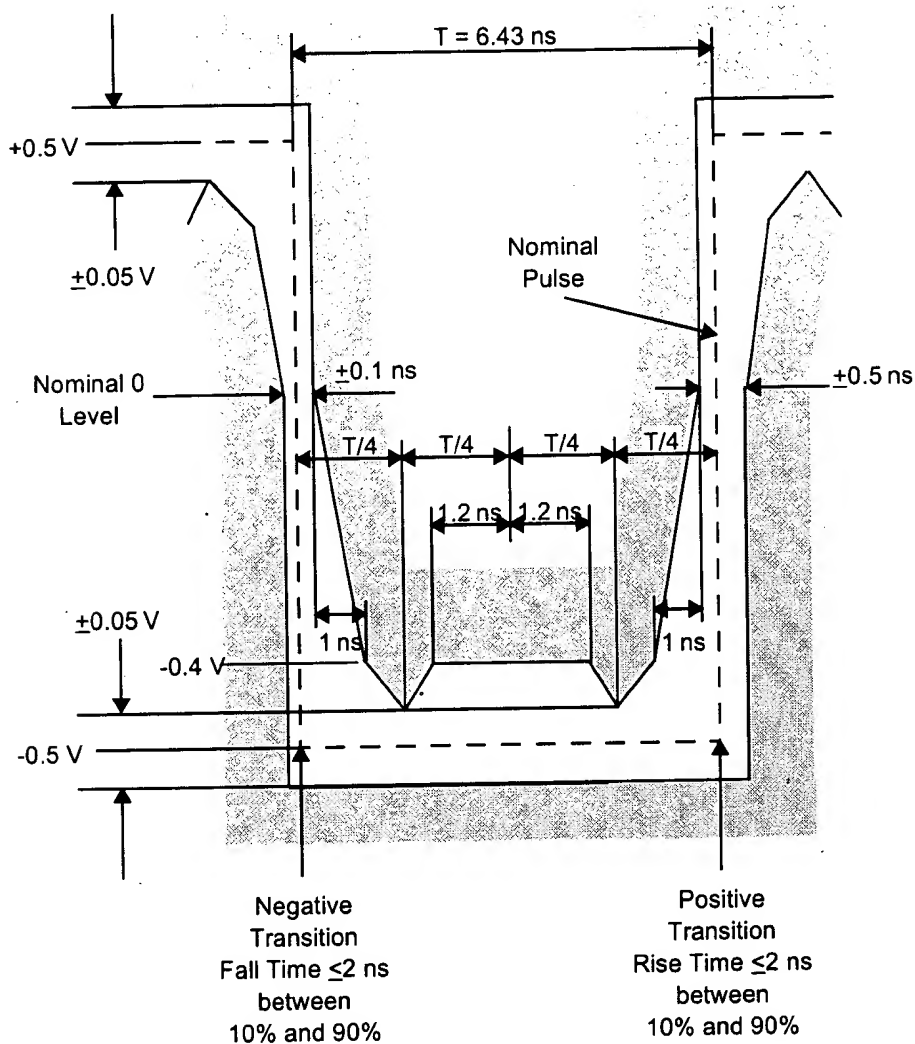
Table 4-13. STS-3 Electrical Interface Criteria

Line Rate	155.520 Mb/s (synchronized to the NE clock; see Section 5.4)
Line Code	Coded Mark Inversion (CMI)
Impedance	A resistive test load of 75 Ω ($\pm 5\%$) is used at the interface for the evaluation of pulse shape and the electrical parameters specified below.
Pulse Shape	An eye diagram of the interface signal shall fit within the limits of the eye diagram mask specification in Figure 4-14. The eye diagram is obtained by using a triggering signal at twice the bit rate (i.e., at 311.040 MHz).
Power Level	A wideband power level measured using a low-pass filter with a flat pass band and a 3-dB cutoff frequency of at least 311.04 MHz shall be between -2.5 dBm and +4.3 dBm.
DC Offset	There shall be no dc power flow across the interface.



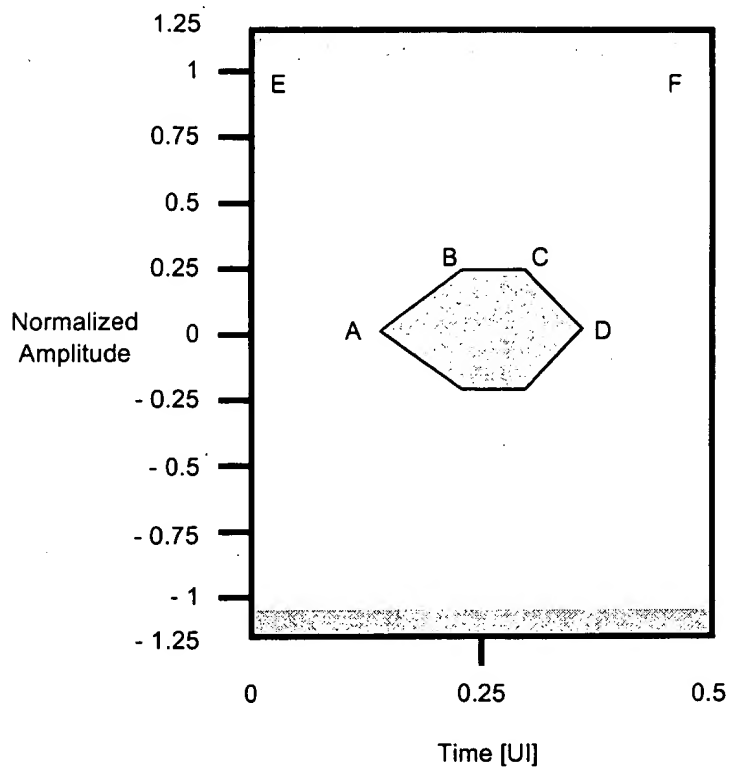
Note 1. The mask does not include the over/undershoot tolerance of 10%.
Note 2. The nominal zero level can be adjusted by $\pm 0.05 \text{ V}$ to meet the limits of the mask.

Figure 4-12. STS-3 Transmitter Pulse Mask Corresponding to a Binary Zero



- Note 1. The mask does not include the over/undershoot tolerance of 10%.
Note 2. The nominal zero level can be adjusted by $\pm 0.05 \text{ V}$ to meet the limits of the mask.
Note 3. The mask also applies to the inverse pulse.

Figure 4-13. STS-3 Transmitter Pulse Mask Corresponding to a Binary One



Corner Points		
Point	Time	Amplitude
A	0.125	0.0
B	0.225	0.25
C	0.275	0.25
D	0.35	0.0
E	0.0	1.1
F	0.5	1.1
NOTE — The mask is symmetric about zero amplitude axis.		

Figure 4-14. STS-3 Eye Diagram Mask

5. Network Element Architectural Features

This section describes various architectural features required in SONET NEs, particularly:

- Multiplexing procedures
- Overhead function usage
- Linear Automatic Protection Switching (APS)
- Network synchronization
- Framing
- Jitter and wander performance.

5.1 Multiplex Procedures

5.1.1 Interleaving

An STS-N module either can be created directly (e.g., by the equipment that creates an STS-12 and maps ATM traffic into the STS-12c SPE), or it can be formed by byte-interleaving lower-level modules (e.g., STS-1s and STS-Ms). For STS-Ns that are formed by byte-interleaving lower-level modules, the following requirements are applicable:

- R5-1** [131] Before byte-interleaving to form an STS-N, the transport overhead byte positions of all the constituent STS-1s and STS-Ms shall be frame aligned.

The alignment of the STS-1s and STS-Ms is accomplished by adjusting the STS Payload Pointers to reflect the new relative positions of the STS SPEs.¹

Note that in development of these requirements, it is useful to assume that an NE that forms an STS-N ($N > 3$) will first logically interleave any STS-1 inputs (in sets of three consecutive STS-1s) to form STS-3 modules, and then interleave those STS-3 modules and any other STS-M inputs to form the STS-N. However, there are no requirements concerning the internal architectures of SONET NEs (e.g., an NE could directly interleave

1. In the SONET interface layer model described in Section 3.3, it would be the responsibility of the Line layer to align and interleave the signals that it receives from the Path layer. These criteria are written based on the assumption that the NE uses Path and Line layer signals that contain all of the Transport Overhead byte positions (e.g., that an internal path layer signal carrying an STS-1 SPE has a nominal bit rate of 51.84 Mb/s); however, as discussed in Section 3.3, there are no criteria concerning the internal signals used by an NE. If the NE uses some other type of internal signals, then it must perform an equivalent process so that the structure of the interface signals (i.e., the OC-N or STS-N electrical signals transmitted by the NE) is independent of the type of internal signals used.

STS-1 inputs in the appropriate order to form an STS-N). The important point is the output byte sequence must be as shown in Figure 5-1.

R5-2 [132] To form an STS-3 from STS-1s, three STS-1s shall be interleaved, one byte at a time. The first byte of the STS-3 shall be the A1 byte from the first STS-1, followed sequentially by the A1 byte from the second STS-1, and then the A1 byte from the third STS-1.

R5-3 [133] To form a higher-level STS-N ($N > 3$) from lower-level STS-Ms ($3 \leq M < N$), the STS-Ms shall be interleaved $M/3$ bytes at a time. The output byte sequence shall be as shown in Figure 5-1.

The preceding requirements have several implications concerning the transport of STS-Mc SPEs in OC-N ($M < N$) signals. Since the STS-Mc SPE is a single entity (see Section 3.2.3), the module that contains it is considered an STS-M input to the byte-interleaver (which has an STS-N as its output). Based on the above requirements, the STS-M input must occupy the byte positions corresponding to M consecutive STS-1 inputs, and it must start in a byte position corresponding to an STS-1 number $[(X+1), 1]$, where $0 \leq X \leq (N-M)/3$ (using the two-level numbering scheme shown in Figure 5-1). In addition, to simplify the processing capabilities required of NEs that originate, pass, and terminate concatenated signals, the possible starting positions for an STS-Mc SPE in an STS-N are further limited by the following requirement.

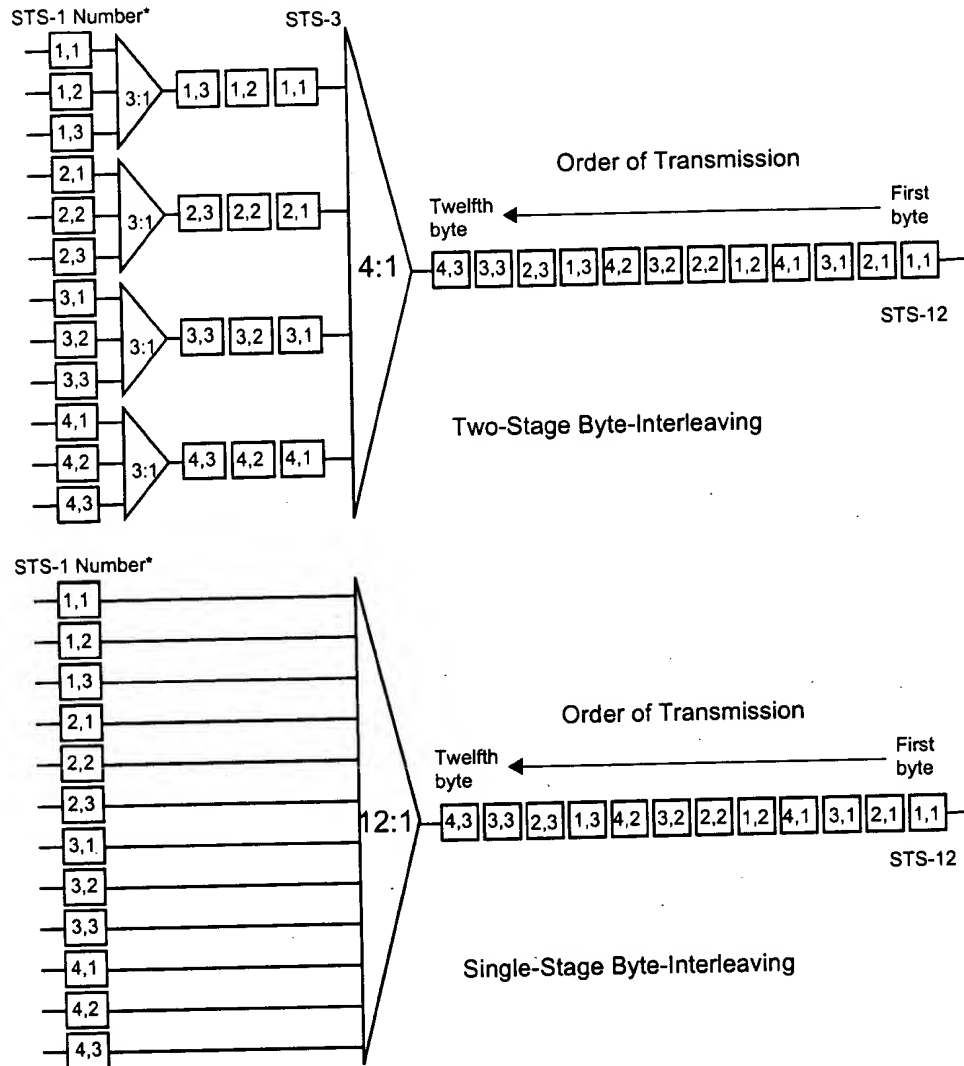
R5-4 [911] Each STS-Mc SPE in an STS-N shall be completely contained in one of Y groups of P STS-1s, where $Y = (N+P)$ and P is the smallest of the three numbers '3', '12', and '48' that satisfies the inequality $M \leq P \leq N$.

The effect of **R5-4** [911] is shown in Table 5-1, which lists the possible starting positions for the various size STS-Mc SPEs that could be contained in an OC-48 signal. Note that although all of the possible values of M from 3 to 48 are included in Table 5-1, many SONET products may support the transport of only certain size STS-Mc SPEs (i.e., STS-3c, STS-12c and STS-48c SPEs). Therefore, an STS-Mc path terminating product's use of other values of M may restrict its applicability in multi-product configurations.

R5-5 [912] A SONET NE that provides the capability to terminate or pass a particular size STS-Mc SPE shall be capable of performing that function on an STS-Mc SPE that starts in any of the allowed starting positions defined in **R5-4** [911] and shown in Table 5-1.

Figure 5-2 illustrates an interleaving example in which nine STS-1s, one STS-3c (carrying an STS-3c SPE), one STS-12c (carrying an STS-12c SPE), and other STS-1 and STS-Mc modules (which are unspecified in the example) are interleaved to form an STS-48. In the example, the term "first 4 bytes of STS-12c 4,1" means the first four bytes that would be transmitted if that STS-12 were converted to an OC-12 signal, rather than multiplexed into

an STS-48. Similarly, “first byte of STS-3c 8,1” means the first byte that would be transmitted if that STS-3 were converted to an STS-3 electrical or OC-3 signal, etc.



*The STS-1s in this figure are numbered using the two-level “STS-3 #/STS-1 #” numbering scheme. This scheme is based on an STS-3 number followed by the number of the STS-1 within the STS-3, and is one of the two schemes that may be used by SONET NEs (see Section 6.1.2). The other scheme that may be used is one in which the STS-1s are numbered “1 to N in order of appearance at the input to the byte-interleaver”. Using that scheme, the order of transmission in this figure would be 1 (first byte), 4, 7, 10, 2, 5, 8, 11, 3, 6, 9, 12 (twelfth byte).

Figure 5-1. Example of Byte-Interleaving Sequence, STS-12

Table 5-1. Possible Starting Positions for an STS-Mc SPE in an OC-48 Signal

STS-1 Number ^a	STS-1 Number ^b	STS-3c SPE	STS-6c SPE	STS-9c SPE	STS-12c SPE	STS-15c SPE	STS-18c SPE	STS-21c to STS-45c SPEs	STS-48c SPE
1,1	1	Y	Y	Y	Y	Y	Y	Y	Y
2,1	4	Y	Y	Y	NO	Y	Y	Y	NO
3,1	7	Y	Y	NO	NO	Y	Y	..	NO
4,1	10	Y	NO	NO	NO	Y	Y	..	NO
5,1	13	Y	Y	Y	Y	Y	Y	..	NO
6,1	16	Y	Y	Y	NO	Y	Y	..	NO
7,1	19	Y	Y	NO	NO	Y	Y	..	NO
8,1	22	Y	NO	NO	NO	Y	Y	..	NO
9,1	25	Y	Y	Y	Y	Y	Y	..	NO
10,1	28	Y	Y	Y	NO	Y	Y	..	NO
11,1	31	Y	Y	NO	NO	Y	Y	NO	NO
12,1	34	Y	NO	NO	NO	Y	NO	NO	NO
13,1	37	Y	Y	Y	Y	NO	NO	NO	NO
14,1	40	Y	Y	Y	NO	NO	NO	NO	NO
15,1	43	Y	Y	NO	NO	NO	NO	NO	NO
16,1	46	Y	NO	NO	NO	NO	NO	NO	NO

Notes:

- The STS-1 numbers shown in the first column are the two-level "STS-3 #/STS-1 #" numbering scheme that SONET NEs may use (see Section 6.1.2). Note that an STS-Mc SPE cannot start in an STS-1 numbered (1,2), (1,3), (2,2), (2,3), ..., or (16,3).
- The STS-1 numbers shown in the second column are the single-level "1 to N in order of appearance at the input to the byte-interleaver" numbering scheme that SONET NEs may use (see Section 6.1.2). Note that an STS-Mc SPE cannot start in an STS-1 numbered 2, 3, 5, 6, ..., or 48.

Y = STS-Mc SPE can start in that STS-1

NO = STS-Mc SPE cannot start in that STS-1

.. = Y or NO, depending on the particular value of M

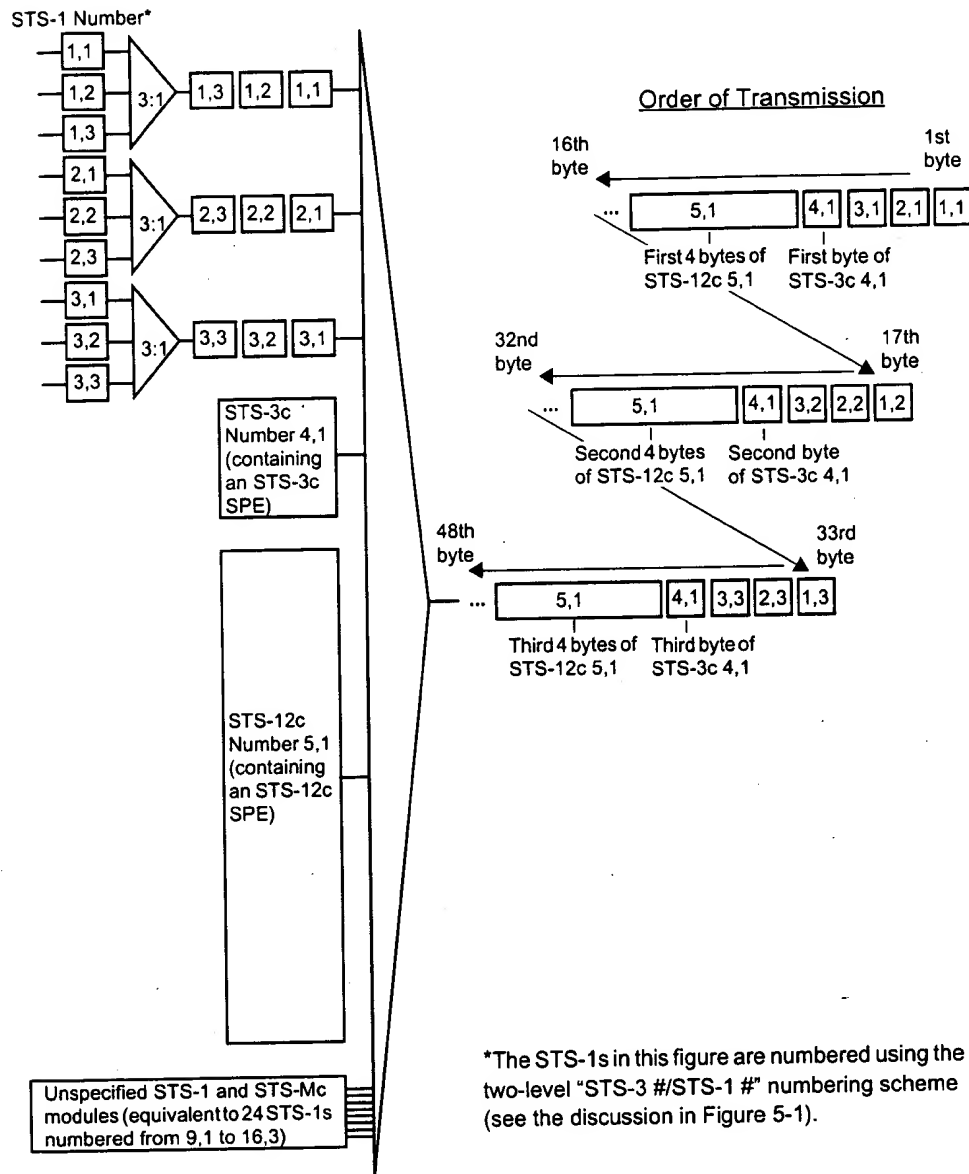


Figure 5-2. Byte-Interleaving Example, Multiple Level Inputs

5.1.2 Concatenation

In all of the uses of STS-Nc SPEs currently defined in SONET, the STS-Nc SPE contains a single payload mapping (e.g., the ATM mapping into an STS-12c SPE described in Section 3.4.3.2.1), and therefore it is created in a single STS PTE. Section 3.2.3 contains criteria related to the structure of STS-Nc SPEs, Section 3.4.3 describes the mappings of Super Rate payloads into STS-Nc SPEs, and Section 3.5.1.4 describes the use of STS Payload Pointers to identify the STS-1s that make up an STS-Nc. If future uses are defined that require an STS-Nc SPE to carry multiple payloads mapped into STS-1 or STS-Mc ($M < N$) SPEs, then the necessary additional criteria will be added to this section.²

5.1.3 Scrambling

SONET optical interface signals use binary line coding, and therefore must be scrambled to assure an adequate number of transitions (zeros to ones, and ones to zeros) for such purposes as line rate clock recovery at the receiver. SONET electrical interface signals use line codes that assure adequate transitions (i.e., B3ZS and CMI, see Section 4.4); however, they are also scrambled for consistency between the electrical and optical interfaces. In both cases, the scrambler used is a frame synchronous scrambler that can be applied identically at the transmitter and the receiver.

R5-6 [134] SONET interface signals shall be scrambled (i.e., scrambled at the transmitter and descrambled at the receiver) using a frame synchronous scrambler of sequence length 127, operating at the line rate.

... The generating polynomial for the scrambler shall be $1+x^6+x^7$.

... The scrambler shall be reset to '111111' on the most-significant bit of the byte following the Z0 byte in the Nth STS-1 (i.e., the byte following the last Z0 byte). That bit and all subsequent bits to be scrambled shall be added, modulo 2, to the output from the x^7 position of the scrambler, as shown in Figure 5-3. The scrambler shall run continuously from that bit on throughout the remainder of the STS-N frame.

Note that the framing bytes (A1 and A2), the Section Trace byte (J0), and the Section Growth (Z0) bytes are not scrambled.

2. Such a use has been defined in SDH, but no equivalent use is defined in SONET. In the SDH case, a single Virtual Container-4 (VC-4, the SDH equivalent to an STS-3c SPE) is used to transport three independent Tributary Unit Group-3s (TUG-3s). Each TUG-3 is equivalent to an STS-1 SPE with the fixed stuff columns removed and a new column containing a payload pointer added. This additional pointer is used to identify the start of the VC-3, which in turn can contain (for example) an asynchronously mapped DS3 signal.

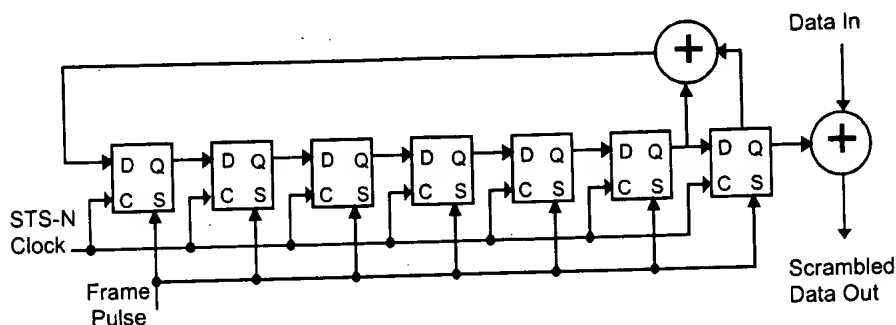


Figure 5-3. Frame Synchronous Scrambler (Functional Diagram)

5.1.4 An Example of STS-1 and OC-N Signal Composition

Figure 5-4 shows one possible set of stages in the formation of an STS-1 and an OC-N signal. The NE in this example has STE, LTE, and PTE functionality, and is also capable of passing nonterminated STS SPEs through. This section merely illustrates the functions, and is not intended to be a description of a required implementation.

The content of a terminating STS-1 SPE is formed from a payload mapping (e.g., a DS3 signal into an STS-1 SPE). Included within the STS-1 SPE are the nine bytes that make up the STS POH. Super Rate payloads can be mapped into an STS-Nc SPE, which also includes nine bytes of STS POH. The BIP-8 error check byte is calculated over the entire STS SPE and the result is placed in the B3 byte of the following STS SPE. If the path termination is not equipped or provisioned, then an "Unequipped" STS SPE is indicated by inserting an all-zeros pattern.

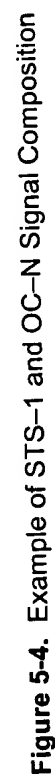
If the contents of a nonterminating STS SPE are lost (e.g., as the result of a failure at an OC-M input to an OC-M to OC-N multiplex), then an all-ones pattern (for STS Path AIS) is inserted into the entire STS SPE and its associated H1 through H3 bytes.

The next stage involves frame aligning the STS-1s and STS-Ms, and the addition (or overwriting) of the Line Overhead. The BIP-8 (B2), Payload Pointer and Pointer Action (H1, H2, H3), and Growth (Z1 and Z2) bytes are present in the Line Overhead of each (or in the case of Z1 and Z2, all but one) STS-1 in an STS-N. Thus, Line level BIP-8 error checking and STS Payload Pointer processing are actually performed on each individual STS-1, irrespective of whether that STS-1 is part of an STS-M carrying an STS-Mc SPE.

The remaining bytes of the Line Overhead are defined only for a single STS-1 in an OC-N signal (e.g., the S1 byte in the first STS-1, the M1 byte in the third STS-1). An all-zeros pattern is sent over the undefined bytes within the other STS-1s. The STS-N is obtained by byte-interleaving the STS-1s and STS-Ms in the proper order.

The final part of the overhead, the Section Overhead, is added (or overwritten) next. The Section BIP-8 (B1), Orderwire (E1), User Channel (F1), and Data Communications Channel (D1 through D3) bytes are present only in the first STS-1 of an OC-N signal, and are all added before scrambling. An all-zeros pattern is added in the corresponding undefined bytes in the second through Nth STS-1s. The STS-N is scrambled, and then the framing (A1, A2) bytes and Section Trace (J0) or Growth (Z0) bytes are added (or overwritten) for each STS-1. The final operation is the calculation of the Section BIP-8 over each 125 μ s frame of the scrambled STS-N signal. The result is placed into the B1 byte in the following frame.

The STS-N signal is then converted into optical pulses for transmission over the fiber.



5.2 Overhead Function Usage

5.2.1 Generating and Processing Overhead

Table 5-2 summarizes the criteria for generating and processing overhead bytes for SONET NEs, and references the primary sections in this GR that contain the applicable criteria.³ Generating overhead refers to the creation of the overhead bits or bytes for the transmitted SONET signal, while processing refers to the interpretation and, if required, the further derivation of the data from the information in the overhead bits and bytes of the received signal. In general, the criteria apply only to the layers that are terminated by a particular NE. For example, the STE, LTE, STS PTE and VT PTE criteria would all be applicable to two ADMs with DS1 low-speed interfaces, which generate and terminate an OC-N signal. However, only the STE criteria would apply to an STE regenerator that regenerates that same OC-N signal. Note however, that for some applications (e.g., intermediate path PM, see Section 6.2.2.9) it may be necessary for an NE to process overhead in a layer that it does not terminate.

The status "R" in Table 5-2 denotes that overhead generation or processing is required, either for all applications or for certain applications (as noted). The status "CR" denotes that the specified use of the overhead may be required in certain applications. The status "O" denotes that the specified use of the overhead is an objective in certain applications. For other applications, or where the status is unmarked, the overhead is either undefined (generation) or unused (processing). See Section 3.2 for the criteria on undefined and unused overhead.

- R5-7** [135] Overhead that is required to be generated shall carry valid data as this GR describes. Processing for the required overhead shall adhere to the criteria contained in this GR or the appropriate NE-specific GRs, TRs, and TAs.

Note that the particular signals for which the various overhead bits and bytes listed in Table 5-2 need to be processed will be affected by the type of APS architecture (if any) that is supported. For example, based on the definition of the linear APS protection switching boundaries in Section 5.3.1, each incoming SONET signal in a linear APS system is separately monitored for several items that are required to be detected on a per-line basis for protection switching and Line Performance Monitoring purposes [e.g., Line BIP (B2 byte) errors, AIS-L (or lower-layer LOS and SEF/LOF) defects, RDI-L defects, REI-L indications]. In addition, the detection of certain of those items on any particular signal (e.g., on the incoming signal on the protection line) is directly reflected in the generation of REI-L and RDI-L indications in the Line overhead on the corresponding upstream signal (e.g., in the outgoing signal on the protection line). On the other hand, the NE may process

3. Overhead bits or bytes for which the only currently applicable criteria are those in Section 3.2 are generally not shown in Table 5-2 (e.g., Z1, Z2, G1 bit 8, Z4, V4, J2, Z6, and Z7 bits 1, 2, 3, 4, and 8).

various other Section and Line overhead bytes such as the orderwire bytes (E1 and E2), the DCC bytes (D1 to D3, and D4 to D12), the STS payload pointer bytes (H1 and H2), and the synchronization status byte (S1) in any of the following:

- All of the incoming signals (but in most cases only use the information from the active line or lines)
- Only the incoming signal or signals from the selected line or lines
- None of the incoming signals (e.g., the orderwire bytes if orderwire is not being used).

Similarly, in a UPSR NE operating at the STS level, the STS payload pointer bytes and certain STS Path overhead bytes are monitored for protection switching and possibly Performance Monitoring purposes. However, since that STS Path is not necessarily terminated within that same NE (in contrast to the linear APS case where the Line layer is always terminated by the NE that performs the switching), the results of that monitoring are not directly reflected in the upstream STS Path signals. Instead, the RDI-P and REI-P indications in the upstream STS Path signals are based on defects and errors that are detected on the STS Path that is eventually terminated.

In addition to processing overhead bits and bytes in nominal-rate SONET signals, it is also important for a SONET NE to be able to accept signals that are off frequency. In the following requirement, to "be capable of receiving and processing" a signal is intended to mean that the NE meets the applicable physical criteria (e.g., receiver sensitivity, jitter tolerance), and can perform the SONET-related functions that it normally performs (e.g., regenerate the signal, process the SONET overhead for the layers that it terminates, process the pointers for various paths, terminate the appropriate paths). It is not intended to mean that non-SONET payload signals dropped by the NE are required to be error-free when the incoming SONET signal is greater than 4.6 ppm off frequency, or that the signal must be accepted as a timing reference signal. Separate criteria (e.g., **O5-238 [349]**, and the **GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria***, criteria referenced in Section 5.4.6.1 of this document) address those issues. Finally, the requirement is applicable when the NE's internal clock is operating anywhere within its own free-run accuracy specifications. For example, an NE with an SMC that is free-running at -20 ppm must be able to receive and process an incoming SONET signal with a bit-rate that is +20 ppm off frequency.

R5-8

[1013] A SONET NE shall be capable of receiving and processing incoming SONET signals with bit-rates that are, at a minimum, anywhere in the range of ± 20 ppm off frequency from the nominal bit-rates for those signals.

Table 5-2. SONET Overhead Generating and Processing Criteria

	OH Bytes	Function or Definition	Generation Criteria		Processing Criteria	
			Status	Section	Status	Section
S T E	A1, A2	Framing	R	3.3.2.1	R ¹	5.5
	J0	Section Trace/Fixed Pattern	O ² (Trace)/ R(Fixed)	3.3.2.1		
	Z0	Section Growth (Fixed Pattern)	R	3.3.2.1		
	B1	Section BIP-8	R ²	3.3.2.1	R ^{1,2,3}	6.2.2.3
	E1	Local Orderwire	CR ^{2,4}	5.2.2	CR ^{2,4}	5.2.2
	F1	Section User Channel	CR ^{2,4}	3.3.2.1	CR ^{2,4}	5.2.3
	D1, D2, D3	Section DCC	R ^{2,4,5}	8.3.1.3	R ^{2,4,5}	8.3.1.3
L T E	H1, H2, H3	STS Payload Pointer and Action	R ⁶	3.5.1	R ⁶	3.5.1
	B2	Line BIP-8	R	3.3.2.2	R	6.2.2.4
	K1, K2	APS Channel	R ⁷	5.3.5	R ⁷	5.3.5
	K2 bits 6-8	RDI-L AIS-L Monitoring	R	6.2.1.3.1	R R	6.2.1.3.1 6.2.1.2.1
	D4-D12	Line DCC	CR ^{2,4,5}	8.3.1.3	CR ^{2,4,5}	8.3.1.3
	S1	Synchronization Status	R ⁴	5.4.2	R ⁴	5.4.2
	M0/M1	REI-L	R	3.3.2.2	R	6.2.2.4
	E2	Express Orderwire	CR ^{2,4}	5.2.2	CR ^{2,4}	5.2.2
S T S P T E	H1, H2, H3	STS Payload Pointer and Action AIS-P Monitoring (H1 and H2 only)	R ⁶	3.5.1	R ⁶ R	3.5.1 6.2.1.2.2
	J1	STS Path Trace and TIM-P	R	6.2.1.1.9 6.2.3.2.3	R	6.2.1.1.9 6.2.3.2.3
	B3	STS Path BIP-8	R	3.3.2.3	R	6.2.2.5
	C2	STS Path Signal Label PDI-P Monitoring	R	3.3.2.3	R CR	6.2.1.1.8 6.2.1.4.1
	G1 bits 1-4	REI-P	R	3.3.2.3	R	6.2.2.5
	G1 bit 5	One-bit RDI-P	R	6.2.1.3.2	R	6.2.1.3.2
	G1 bits 5-7	ERDI-P	O	6.2.1.3.2	O	6.2.1.3.2
	F2	Path User Channel ⁸	CR	3.3.2.3	CR	5.2.3
	H4	Indicator	R ⁹	3.4	R ⁹	3.4
	Z3	Growth ⁸				
	Z5	Tandem Connection and Path Data		3.3.2.3		
	V1, V2, V3	VT Payload Pointer and Action	R ¹⁰	3.5.2	R ¹⁰	3.5.2
V T P T E	V1, V2, V3	VT Payload Pointer and Action AIS-V Monitoring (V1 and V2 only)	R ¹⁰	3.5.2	R ¹⁰ R	3.5.2 6.2.1.2.3
	V5 bits 1, 2	VT Path BIP-2	R	3.3.3	R	6.2.2.6
	V5 bit 3	REI-V	R	3.3.3	R	6.2.2.6
	V5 bit 4	RFI-V	R ¹¹	6.2.1.3.3	R ¹¹	6.2.1.3.3
	V5 bits 5-7	VT Path Signal Label	R	3.3.3	R	6.2.1.1.8
	V5 bit 8	RDI-V	R	6.2.1.3.3	R	6.2.1.3.3
	Z7 bits 5-7	ERDI-V	O	6.2.1.3.3	O	6.2.1.3.3

See Notes on following page

Notes:

1. Also required for physical layer regenerators (see TR-NWT-000917).
2. Not required, or not applicable, for drop-side signals (see Section 3.3.2 for the definitions of line-side and drop-side signals).
3. Required for regenerators. Conditionally required for the NEs with both STE and LTE functionality in applications with regenerators. Not required in applications with no regenerators.
4. Not required, or not applicable, for lines 2 through n of a 1:n protected system.
5. Not required, or not applicable, for B-ISDN UNI applications.
6. Not required for nonterminated STSs that are dropped from one SONET signal to a lower-speed SONET signal, if all of the STSs in the low-speed signal are from the same high-speed signal and the low-speed signal uses through-timing such that there are no changes in the offsets between the pointers and the first bytes of the STS SPEs (i.e., the incoming pointers are passed through unchanged from the high-speed signal to the low-speed signal). STS Payload Pointers for other nonterminated STSs may be processed at the LTE where the incoming SONET signal is terminated, at the LTE where the outgoing SONET signal is generated, or both. STS Payload Pointers for terminated STSs may be processed at the LTE, the STS PTE, or both.
7. Required only for the protection line.
8. The F2 and Z3 bytes are used for payload-specific functions in the DQDB mapping. See Section 3.4.3.1.4.
9. Currently required only for VT-structured STS-1 SPEs, and the DQDB mapping.
10. VT Payload Pointers for nonterminated VTs may be processed at the STS PTE where the incoming SONET signal is terminated, at the STS PTE where the outgoing SONET signal is generated, or both. The VT Payload Pointers for terminated VTs may be processed at the STS PTE, the VT PTE, or both. [Note that the frequency justification method previously discussed in TR-NWT-000496 (which has since been replaced by GR-496-CORE), of using the STS pointer to justify through VTs when an STS-1 is terminated but only some of the VTs are dropped, is not recommended for most applications. That method can result in unnecessary "hits" on added VTs during a recovery from an incoming signal failure.]
11. Required only for the byte-synchronous DS1 mapping.

5.2.2 Orderwire (OW)

Craftspeople use the OW channels for voice communications between different sites during failures in the network, or during coordinated maintenance activities such as provisioning

and testing of new services. Two channels are allocated for OW in the SONET overhead. These channels are:

- Local Orderwire (LOW), which is carried in the E1 byte of the Section Overhead
- Express Orderwire (EOW), which is carried in the E2 byte of the Line Overhead.

Both of these channels have bit rates of 64 kb/s, and are located in the overhead of the first STS-1 in an OC-N signal. This section contains the OW criteria applicable to a SONET NE.

CR5-9 [136] SONET NEs with line-side interfaces may be required to provide OW functionality.

Support of OW at drop-side interfaces is not required or expected.

The primary distinction between the LOW and the EOW is that the LOW can be accessed at STE (i.e., it can be generated and processed at STE regenerators, and at all NEs that also have LTE functionality), while the EOW can be accessed only at LTE. Therefore, for systems with no regenerators, support of both the LOW and the EOW is not necessary. To promote OW interoperability between NEs that only support one of the OW channels, the following objective is applicable:

O5-10 [137] If only a single OW channel is supported, it should be the LOW.

If an NE does not support either the LOW or the EOW, then either the E1 byte or the E2 byte (respectively) is considered undefined, and the criteria in Section 3.2 are applicable.⁴

5.2.2.1 OW Access

The following OW access criteria apply if OW functionality is supported.

O5-11 [138] Access to the OW circuit should be through a 4-wire analog interface at 0 dBm. The input impedance should be 600 Ω ($\pm 5\%$), and the speech encoding should be μ -law Pulse Code Modulation (PCM).

R5-12 [139] If a 4-wire analog interface is not provided, either a 2-wire analog interface [0 dBm, 900 Ω ($\pm 5\%$), μ -law PCM], or a digital interface shall be provided.

Whether the digital encoding is performed by the SONET NE (as is the case for 4-wire or 2-wire analog interfaces), or by external equipment (digital interfaces), the 8-bit PCM sample must be placed in the OC-N signal as follows:

4. It is also acceptable for the NE to insert the "quiet" PCM code (see Section 5.2.2.1) on the unsupported E1 or E2 bytes.

- R5-13** [140] The 8-bit PCM sample shall be synchronized to the STS-N frame, and the PCM bits shall be assigned to the corresponding bits of the appropriate E1 or E2 byte (see Figure 3-2).

The following requirement is applicable for the OW channel or channels that are supported:

- R5-14** [141] The NE shall have the capability to generate the "quiet" PCM code (01111111) on its supported OW channels.

5.2.2.2 OW System Communication

The LOW is intended to provide voice communications between any two NEs with STE functionality along a SONET line (including the end NEs that have both STE and LTE functionality). This means that all STE must be able to access the E1 bytes. It also means that the NEs that provide only STE functionality (e.g., STE regenerators) must be able to pass⁵ the LOW channel, to allow voice communications between non-consecutive sites on the line.

- R5-15** [142] An NE with STE functionality but no LTE functionality shall provide the capability to pass the incoming LOW channel through to the outgoing LOW channel on the same line.

The EOW is intended to provide voice communications between the LTE on a SONET line without any dependency on the LOW functionality of intermediate STE. This means the LTE at each end of the SONET line must be able to access the E2 bytes to allow end-to-end communication. In addition, in some applications it may be useful for an NE that terminates multiple SONET optical lines to provide the capability to interconnect the EOW channels on different lines.

- CR5-16** [143] A SONET ADM may be required to provide the capability to pass the EOW circuit between the high-speed OC-N signals.
- CR5-17** [144] A SONET NE that terminates multiple SONET optical lines may be required to provide the capability to pass the EOW circuit between any two of those lines.

In some applications, it may be useful for the OW channel to be protected.

5. Typically when the word "pass" is used with respect to a digital signal, it is intended to mean that the incoming and outgoing bit-streams in the same direction of transmission are identical. However, when "pass" is used in this section, it would also be acceptable for the NE to perform back-to-back digital-to-analog and analog-to-digital conversions, and to provide circuitry that would allow multi-point communications. For example, a regenerator could provide party-line access in which the local input is added (in the analog domain) to the "passed" OW channel in each direction, and the received OW channels from both directions are added to become the local output.

CR5-18 [145] A SONET NE may be required to support OW channel protection.

R5-19 [146] If OW channel protection is supported, then the OW protection scheme shall be the overhead protection scheme in which overhead channels are protected along with the traffic (see Section 8.3.1.3).

Note that using this protection scheme, in a configuration where two NEs supporting linear APS are connected via diversely-routed working and protection lines that include one or more regenerators, communications between an end NE and a regenerator site may not be possible. Specifically, if the end NE is selecting traffic from the working line then communications with a regenerator site that supports only the protection line will not be possible, while if the end NE is selecting traffic from the protection line then communications with a site that supports only the working line will not be possible.

5.2.2.3 OW Operations

Monitoring and control capabilities (e.g., selective signaling, ringing capabilities, call-out functions, audible and visual call indications, OW channel selection) to assist users in establishing OW connections are for further study.

5.2.3 User Channels

As discussed in Section 3.3.2, overhead bytes are allocated in the SONET signal for use by the network provider. The following criteria apply to the use of those user channels:

CR5-20 [147] An NE with STE functionality may be required to allow the user to access the Section User Channel (i.e., the F1 byte) in line-side signals.

CR5-21 [148] An NE that provides only STE functionality (e.g., an STE regenerator) may be required to be capable of passing the incoming F1 byte through to the outgoing signal on the same line.

Support of the Section User Channel is not required for drop-side signals.

CR5-22 [149] An NE with STS PTE functionality may be required to allow the user to access the Path User Channel (i.e., the F2 byte).

5.3 Automatic Protection Switching

Several Automatic Protection Switching (APS) schemes have been defined for SONET signals. This section contains criteria related to the scheme known as linear APS. Other Bellcore criteria documents [e.g., GR-1230-CORE, *SONET Bidirectional Line-Switched*

Ring Equipment Generic Criteria, and GR-1400-CORE, *SONET Dual-Fed Unidirectional Path Switched Ring (UPSR) Equipment Generic Criteria*] contain criteria applicable to systems that use other protection switching schemes. In addition, Section 5.6 of GR-499-CORE contains circuit pack protection switching criteria that are applicable to SONET NEs that provide redundant or protection circuit packs.

- CR5-23** [150] SONET LTE that terminates optical lines may be required to provide linear APS.

Support of linear APS at STS-N electrical interfaces is not expected or required.

- R5-24** [151] If linear APS is provided, the SONET NE shall provide the capability for the user to disable the feature on a per-interface basis.

The capability to disable the linear APS feature is needed for interworking with SONET NEs that do not support it.

Linear APS, and in particular, the protocol for the APS channel, is standardized to allow interworking between SONET LTE from different suppliers. The criteria in this section apply to all SONET LTE that provides linear APS. To facilitate the coordinated revision of all of the SONET linear APS criteria, those criteria that previously appeared in TR-NWT-000499 are being repeated here, with the appropriate revisions. Those criteria do not appear in GR-499-CORE.

5.3.1 Protection Switching Boundaries

Linear APS is defined to provide protection at the line layer (see Section 3.3). Therefore, all of the STS SPEs carried in an OC-N signal are protected together (i.e., if a switch occurs, all of the STS SPEs are switched simultaneously).

- R5-25** [152] Protection shall cover the multiplexer/optics units from the point, at or before, where the Line Overhead is inserted (the head end) to the point, at or beyond, where it is terminated (the tail end).

5.3.2 Linear APS Architectures

Two linear APS architectures are defined in the following sections. These are the 1+1 architecture and the 1:n architecture. In addition, the special case of the 1:n architecture where n is equal to 1 is also discussed. Several criteria are given that apply only to the 1:1 case, to allow interworking with LTE using the 1+1 architecture.

5.3.2.1 1+1 Architecture

The 1+1 architecture is defined as follows:

An architecture in which the head-end signal is continuously bridged (at the electrical level) to working and protection equipment so that the same payloads are transmitted identically to the tail-end working and protection equipment. At the tail end, the working and protection OC-N signals are monitored independently and identically for failures. The receiving equipment chooses either the working or the protection signal as the one from which to select the traffic, based on the switch initiation criteria contained in Section 5.3.3. Because of the continuous head-end bridge, the 1+1 architecture does not allow an unprotected extra traffic channel to be provided.

The 1+1 architecture could be supported as either a user-provisionable option (e.g., with the 1:1 case of the 1:n architecture as the other option) or as the only supported architecture. If it is the only supported architecture, it may still be upgradable to the 1:n architecture.

A system using the 1+1 architecture operates, as a default, in a unidirectional mode. In this mode, the switching is complete when a channel in the failed direction is switched to the protection line. Although head-end to tail-end signaling is not needed, the APS channel (which is carried in the K1 and K2 bytes of the signal on the protection line) is still used to indicate the local switch actions and the mode of operation. A bidirectional switching mode may also be provided as a user-provisionable option. In this mode, a channel is switched to the protection line in both directions. Switching of only one direction is not allowed. Head-end to tail-end signaling is accomplished using the APS channel.

A 1+1 system also uses, as a default, nonrevertive switching. In nonrevertive switching, a switch to the protection line is maintained even after the working line has recovered from the failure that caused the switch, or the manual switch command is cleared. Revertive switching may also be provided as a user-provisionable option. In revertive switching, the traffic is switched back to the working line when the working line has recovered from the failure or the manual command is cleared.

CR5-26 [153] The LTE may be required to support the 1+1 architecture.

R5-27 [154] If the 1+1 architecture is provided, the following apply:

- ... • The unidirectional mode shall be provided.
- ... • If bidirectional switching is provided (see **CR5-28** [155]), the mode shall be user-provisionable, with a default mode of unidirectional.
- ... • If bidirectional switching is provided, the switching operation shall be bidirectional only if the LTE at both ends is provisioned to operate in the bidirectional mode. Otherwise, the LTE shall operate in the unidirectional mode. (The K2 byte is used to determine the mode of the far-end LTE as described in Section 5.3.5.2.)

- ... • Nonrevertive switching shall be provided.
 - ... • If revertive switching is provided (see **CR5-29 [156]**), the choice of nonrevertive or revertive shall be user-provisionable, with a default of nonrevertive.
- CR5-28 [155]** If the 1+1 architecture is provided, the bidirectional mode may be required to be provided.
- CR5-29 [156]** If the 1+1 architecture is provided, revertive switching may be required to be provided.
- CR5-30 [157]** 1+1-LTE may be required to be upgradable to the 1:n architecture.

5.3.2.2 1:n Architecture

The 1:n architecture is defined as follows:

An architecture in which any of the n working channels can be bridged to a single protection line. Permissible values of n are from 1 to 14. Head-end to tail-end signaling is accomplished by using the APS channel. Because the head end is switchable, the protection line can be used to carry an extra traffic channel.

The 1:n architecture could be supported as either a user-provisionable option (i.e., with the 1+1 architecture as the other user-provisionable option) or as the only supported architecture (although see Section 5.3.2.3 for criteria related to the 1:1 case).

In a 1:n system, all switching is revertive, and both unidirectional and bidirectional switching modes are provided. The mode of operation is user-provisionable, with bidirectional as the default.

- CR5-31 [158]** LTE may be required to support the 1:n architecture.
- CR5-32 [159]** LTE supporting the 1:n architecture may be required to support values of n greater than 1.
- R5-33 [160]** If the 1:n architecture is provided, the following apply:
- ... • All switching shall be revertive.
 - ... • The unidirectional mode shall be provided.
 - ... • The bidirectional mode shall be provided.
 - ... • The mode shall be user-provisionable, with a default mode of bidirectional.

- ...
- The switching operation shall be unidirectional only if the LTE at both ends is provisioned to operate in the unidirectional mode. Otherwise, the LTE shall operate in the bidirectional mode. (The K2 byte is used to determine the provisioned mode of the far-end LTE as described in Section 5.3.5.2.)

O5-34 [161] If the 1:n architecture is provided, the LTE should provide the capability to transport extra traffic on the protection line when it is not being used for protection.

R5-35 [162] If the capability to transport extra traffic is provided, the SONET NE shall provide the capability for the user to disable that feature for interworking with other SONET NEs that do not support it.

5.3.2.3 1:1 Case of the 1:n Architecture

The 1:1 case is a subset of the 1:n architecture, with n equal to 1. In addition to the criteria for the 1:n architecture, the following criteria are specific to the 1:1 case.

CR5-36 [163] LTE that supports only the 1:1 case of the 1:n architecture may be required to be upgradable to support values of $n > 1$.

The following requirement allows 1:1 LTE to interwork with LTE using the 1+1 architecture.

R5-37 [164] 1:1 LTE (which indicates the 1:n architecture on bit 5 of the transmitted K2 byte) shall operate as 1+1 LTE if the far-end LTE indicates that it is 1+1 LTE, as detected on the received K2 byte. The 1:1 LTE shall continue to indicate the 1:n architecture on the transmitted K2 byte unless it is reprovisioned by the user to 1+1. It shall also continue to indicate its provisioned unidirectional/bidirectional switching mode; however, it shall meet the criteria in Section 5.3.2.1 (instead of the criteria in Section 5.3.2.2) concerning which of those modes to actually operate in.

Although nonrevertive switching is the default for the 1+1 architecture, it is not essential to the operation of a system for both ends to use nonrevertive switching (or for both to use revertive), and therefore the type of switching used locally is not indicated to the far end. Therefore, 1:1 LTE can use either revertive or nonrevertive switching when it is operating as 1+1 LTE.

There are currently no criteria concerning the time at which 1:1 LTE must change its operation to 1+1. The time at which it declares an APS Mode Mismatch failure (see Section 6.2.1.1.6.C) would be an appropriate time.

5.3.3 Switch Initiation and Completion Criteria

5.3.3.1 Switch Initiation Criteria

The following two automatic switch initiation criteria are defined for linear APS.

1. Signal Fail (SF): A "hard failure" condition detected on the incoming OC-N signal.

R5-38 [165] Loss of Signal, Loss of Frame and AIS-L defects (see Section 6.2.1), and a Line BER exceeding 10^{-3} on an incoming OC-N shall be detected as SF conditions on that line.

Other protectable hard failures (e.g., a stuck bit) may also be detected as SF conditions.

CR5-39 [166] The BER threshold for an SF condition may be required to be user-provisionable⁶ over the range of 10^{-3} to 10^{-5} .

2. Signal Degrade (SD): A "soft failure" condition resulting from the Line BER exceeding a pre-selected threshold.

R5-40 [167] A BER exceeding the SD threshold on an incoming OC-N shall be detected as an SD condition on that line.

R5-41 [168] The BER threshold for an SD condition shall be user-provisionable over the range of 10^{-5} to 10^{-9} .

5.3.3.2 Switch Initiation Time

The switch initiation time is the time that it takes LTE to detect an SF or SD condition and initiate a switch (if appropriate).

R5-42 [169] For SF conditions caused by LOS, LOF, or AIS-L defects, the switch initiation time shall be 10 ms or less.

O5-43 [170] For SF conditions caused by LOS, LOF, or AIS-L defects, the switch initiation time should be 8 ms or less.

SF and SD conditions based on the BER criteria are detected when the BER of a line exceeds a threshold. The BER of a line is derived from the sum of Line BIP-8 violation counts of the individual STS-1s in the STS-N. The time allowed to detect that the BER is greater than or equal to a threshold (and to initiate a switch, if appropriate) depends on the actual BER (i.e., it does not depend on the threshold).

6. For both SF and SD thresholds, providing thresholds at 1×10^{-n} , where n is an integer, is sufficient.

In general, the occurrence of errors on an optical signal is expected to be a random process. Therefore, for any particular actual BER, there will be some probability that the number of errors detected in a particular time period will be less than the threshold for detecting an SD or SF condition, and a corresponding probability that the number of errors will be greater than or equal to the threshold. The intent of the following criteria (which were written assuming a Poisson distribution of errors) is to promote the development of BER detection algorithms that meet the goals listed below:

- When the actual BER is greater than or equal to an SD or SF threshold, it is quickly detected as such, and a switch (if appropriate) is initiated.
- The probability of detecting a threshold crossing when the actual BER is below the threshold is very small and tolerant to burst errors.

R5-44 [171] For SF and SD conditions based on BER, the switch initiation time shall be below the "maximum" curve in Figure 5-5 (assuming the actual BER is greater than or equal to the threshold).⁷

O5-45 [172] For SF and SD conditions based on BER, the probability that the switch initiation time will be less than the "objective" curve in Figure 5-5 (for the particular rate OC-N signal) should be greater than 0.95 (assuming the actual BER is greater than or equal to the threshold).

Whereas the "maximum" switch initiation time curve shown in Figure 5-5 applies to all OC-N rates, the objective switch initiation time curves depend on the level N. The curves in Figure 5-5 take into account the error detection saturation effect at high BERs.

R5-46 [173v2] For an SF or SD detection threshold of 10^{-n} and an actual BER of $1 \times 10^{-(n+1)}$ or less, the probability that the SD or SF condition will be detected in the "maximum" switch initiation time from Figure 5-5 (for that particular threshold) shall be less than or equal to 10^{-6} .

For example, if the SD threshold is 10^{-5} and the actual BER is 1×10^{-6} , the probability that an SD condition will be detected during any 1-second interval must be less than or equal to 10^{-6} .

In an analysis situation, it is likely that the BER detection algorithm details needed to calculate LTE's conformance to these criteria (using the assumption of a Poisson distribution of errors) will not be available. Therefore, tests will need to be performed where errors are inserted onto an actual signal. Several methods of generating errors are possible, including attenuation of the optical signal and insertion of errors using a SONET test set. If errors are inserted using a SONET test set, their distribution may be either

7. For a random distribution of errors, it is impossible to specify a true "maximum" switch initiation time (i.e., one where the probability of detecting a threshold crossing and initiating a switch within the specified time is 1.0). However, the probability that the switch initiation time will be within the "maximum" time from Figure 5-5 must be very close to 1.0.

periodic or "random" (e.g., generated using some type of random number generation process).

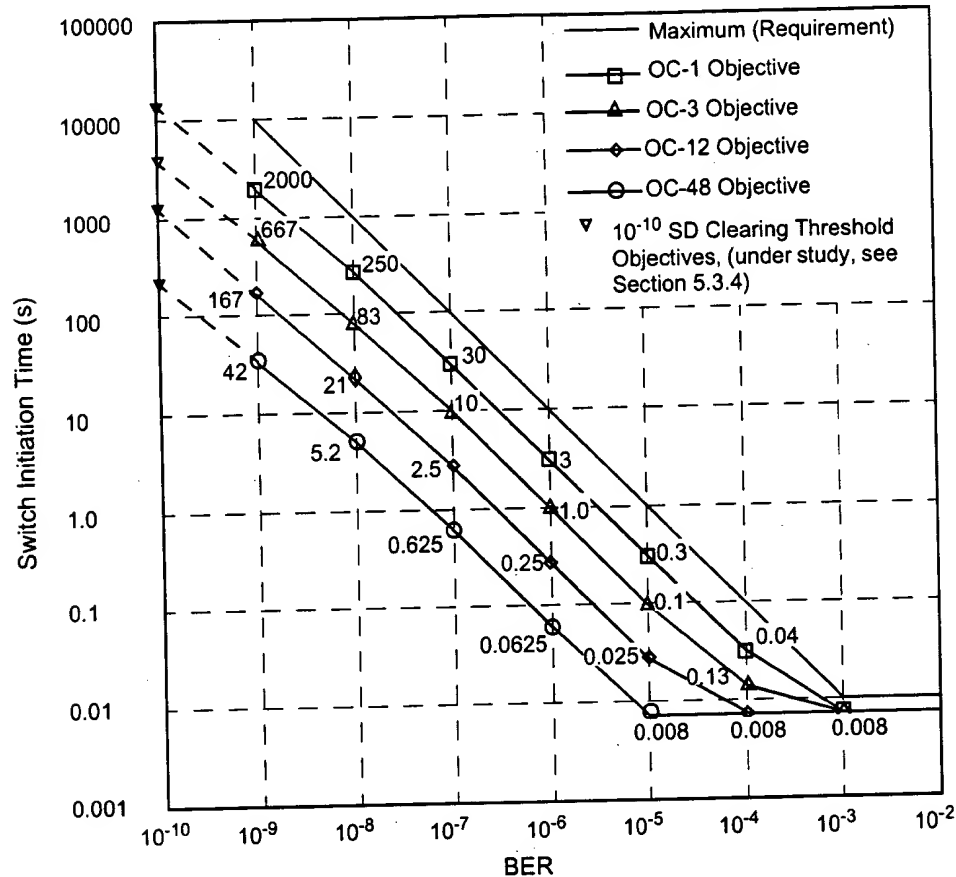


Figure 5-5. Switch Initiation Time Criteria

5.3.3.3 Switch Completion Time

- R5-47** [174] The time to complete a switch, once it is initiated, shall be 50 ms or less.
- R5-48** [175] For bidirectional switching, the LTE at both ends shall complete the switch within the same 50-ms switch completion time, from the time the request is initiated.

For LTE operating in the 1:n architecture or the 1+1 bidirectional mode, the LTE at both ends of the line must perform various actions before a switch can be completed. No criteria exist concerning the times for the individual actions. Therefore, in configurations where the LTE is made by different suppliers, the overall switch completion time may be greater than 50 ms. The need for additional criteria in this area is for further study.

In addition, note that **R5-47 [174]** and **R5-48 [175]** are the applicable requirements for automatically initiated switches caused by degradations or failures, and also for manually initiated switches (e.g., switches initiated by the user from the craft interface). This means that traffic could be interrupted for as long as 50 ms during a manually initiated switch. Some users may find an interruption of that length to be unacceptable, and therefore the following conditional requirement applies for manually initiated switches.

CR5-49 [913] Manually initiated facility protection switches may be required to be error-free.

Note that in many applications, differences in the transmission delays for the working and protection lines may preclude error-free switching. However, it should still be possible for the traffic interruptions caused by a manually initiated switch to be much shorter than the 50 ms switch completion time requirement (i.e., the switch initiation time criteria in Section 5.3.3.2 are not applicable for manually initiated switches, so it should be possible for the NE to take additional time to prepare for the switch, minimizing the time it takes to complete the switch after it is initiated).⁸

Finally, it should be noted that the switch initiation time criteria in Section 5.3.3.2 and the switch completion time requirement in this section are separately applicable. They are not intended to be combined into (for example) a 60 ms total switching time requirement for hard failures. Therefore, if a particular system (e.g., a system that uses a unidirectional switching architecture and semiconductor switching mechanisms) is capable of completing switches much faster than 50 ms, that system is still required to meet the switch initiation time requirements in Section 5.3.3.2 and therefore will have a total switching time for hard failures that is much shorter than 60 ms.

5.3.4 Restoral and Clearing of SD and SF Conditions

For LTE using both revertive and nonrevertive switching, a hysteresis method of clearing SD and SF conditions based on the BER is used. This method specifies an SD or SF clearing threshold ten times lower than the SD or SF threshold.

8. There currently are no criteria in this GR concerning switch initiation times for manually initiated switches. However, other Bellcore requirements documents (e.g., TR-TSY-000824, *OTGR Section 10.1: User System Interface – User System Access*) indicate that an NE must respond to user input within approximately 2 seconds. If a user enters one of the switch commands discussed in Section 5.3.6.1 and that switch is not successfully completed within 2 seconds, it is assumed that the NE will indicate to the user that the command could not be completed. Therefore, the switch initiation time would need to be somewhat less than 2 seconds. The need for explicit criteria in this area is under study.

- R5-50** [176] The clearing threshold for an SD or SF condition based on the BER shall be one-tenth the threshold for declaring the SD or SF.

For 1:n systems, determining that the working line is operating at a BER lower than the clearing threshold needs to be reasonably fast so that the system can restore and free up the protection line for use by other working channels (or the extra traffic channel). In 1+1 systems, the protection line is not needed to carry other channels; however, SD and SF conditions still need to be cleared in a reasonable amount of time to allow the use of the better line when the other line also has an SD or SF condition. On the other hand, it is also important that an existing SD or SF condition not be cleared unless the BER is actually better than the corresponding clearing threshold.

- R5-51** [1046] If an SD or SF condition has been detected and the incoming signal's BER is greater than or equal to that SD or SF threshold, the probability that the LTE will detect that the BER is less than the SD or SF clearing threshold within the "maximum clearing" time listed in Table 5-3 shall be less than or equal to 10^{-6} .
- R5-52** [1047] If an SD or SF condition has been detected and the incoming signal's BER is less than or equal to the SD or SF clearing threshold, the probability that the LTE will detect that the BER is less than that threshold within the "maximum clearing" time shown in Table 5-3 shall be greater than or equal to 0.99.
- O5-53** [177v3] If an SD or SF condition has been detected and the incoming signal's BER is less than or equal to the SD or SF clearing threshold, the probability that the LTE will detect that the BER is less than that threshold within the "objective clearing" time shown in Table 5-3 should be greater than 0.95.

Table 5-3. Clearing Time Criteria for BER-based SF and SD Conditions

SD or SF Threshold	Clearing Threshold	Maximum Clearing Time	Objective Clearing Time / SONET Rate ^a
10^{-3}	10^{-4}	10 ms	None
10^{-4}	10^{-5}	100 ms	10 ms / OC-48 25 ms / OC-12
10^{-5}	10^{-6}	1 s	62.5 ms / OC-48 250 ms / OC-12
10^{-6}	10^{-7}	10 s	625 ms / OC-48 2.5 s / OC-12
10^{-7}	10^{-8}	100 s	5.2 s / OC-48 21 s / OC-12 83 s / OC-3
10^{-8}	10^{-9}	1,000 s	42 s / OC-48 167 s / OC-12 667 s / OC-3
10^{-9}	10^{-10}	10,000 s	340 s / OC-48 1,330 s / OC-12 5,360 s / OC-3

a. Where no objective time is listed for a particular SONET rate, **O5-53 [177v3]** is not applicable.

For consistency with the switch initiation time criteria, the clearing detection times shown above are based on the curves in Figure 5-5. However, in contrast to the switch initiation time criteria, which are a function of the OC-N rate and the actual BER, the clearing time criteria are a function of the OC-N rate and the clearing threshold. For example, for an OC-12 and an SD clearing threshold of 1×10^{-7} , the clearing detection time requirement is 10 seconds and the objective is 2.5 seconds whether the actual BER is 9×10^{-8} or 1×10^{-10} . In this example the SD threshold would be 1×10^{-6} , and the switch initiation time criteria would be 10 seconds (requirement) and 0.25 seconds (objective) for an actual BER of 1×10^{-6} and 0.1 seconds (requirement) and 0.008 seconds (objective) for an actual BER of 1×10^{-4} .

R5-54 [914] Once LTE has detected that the BER is less than the SD or SF clearing threshold, it shall clear the SD or SF condition within an additional 10.5 seconds (although see below for possible exceptions in cases of intermittent SD and SF conditions) assuming the LTE does not detect that the BER is greater than or equal to the SD or SF threshold before that condition is cleared.

Depending on the SONET rate and the particular SD and SF thresholds that LTE is using, **O5-53 [177v3]** could result in extremely short detection times for the purposes of clearing

SD and SF conditions (e.g., for an OC-48 SD threshold of 10^{-5} , the required SD clearing threshold is 10^{-6} and the detection time objective is 62.5 ms). While it is important to clear SD and SF conditions fairly quickly (as discussed above), very short clearing times could result in rapid switches between channels in some situations. For example, oscillations could occur if an SD condition was detected on one line and an intermittent SF condition was detected and cleared repeatedly on another line. Therefore, the following criteria are applicable:

- CR5-55** [915] LTE may be required to be designed to reduce the chance or number of rapid protection switching oscillations that could occur in multiple failure or degradation situations where one or more of the failures or degradations is intermittent.
- R5-56** [916] If LTE is designed to reduce the chance or number of rapid protection switching oscillations, the method used shall be clearly documented.

Two methods that LTE could use to meet **CR5-55** [915] are:

- Delay clearing SD and SF conditions for approximately 10 seconds (as allowed in **R5-54** [914]) after the LTE has detected that the BER of the incoming signal is less than the clearing threshold. The LTE would then clear the condition only if it determines that the SD or SF threshold has not been crossed during the delay period. If the SD or SF threshold is crossed, then the process for clearing the condition is restarted.
- Monitor how often an SF or SD condition is detected and cleared on a line, and "lock on" to that condition if it is detected more than x times in y seconds. After the LTE has locked on to the condition, it would consider that line to be in an SF or SD condition until it determines that the intermittent failure or degradation is gone (e.g., it would clear the SF condition when it has detected that the BER is less than the SF clearing threshold and no hard failures have occurred for z seconds).

The first method discussed above would reduce the chance of rapid oscillations in some situations; however, in other cases it could extend the time that traffic is interrupted.⁹ The second method would not reduce the chance of rapid oscillations; however, it would limit the number of oscillations that could occur.

If a system uses revertive switching, then frequent automatically initiated switches could also (in addition to the multiple failure situations discussed above) occur as the result of an

9. For example, suppose SF conditions have been detected on both working line 1 and the protection line. If the protection line is repaired first and a delay period is used, the NE would not remove the "SF - null channel" request (see Section 5.3.5) for 10 additional seconds. Therefore, it would not be able to insert an "SF - working channel 1" request for those 10 seconds, and the traffic would remain interrupted during the delay period. (Note that if working line 1 was repaired first, the traffic would be restored as soon as the LOS defect on that line was terminated, which would not be affected by whether there is a delay period for clearing SF conditions.)

intermittent failure or degradation on a single working line. To prevent this, a Wait-to-Restore (WTR) period is defined for LTE using revertive switching. After the BER of the working line meets the clearing threshold and the SD or SF condition is cleared, a WTR period is allowed to elapse before restoral. Note that the WTR period is not used after manually initiated switches, or after an SD or SF condition on the protection line clears.

- R5-57** [178v2] For LTE using revertive switching, a WTR period of 5 to 12 minutes shall be provided after the condition that caused an automatically initiated switch to the protection line clears. The length of the WTR period shall be user-provisionable on a per-protection line (or per-protection group) basis.¹⁰

5.3.5 APS Channel Protocol

This section presents the protocol for the APS channel, which is carried in the K1 and K2 bytes in the signal on the protection line. The APS controllers at the LTE terminating SONET lines use the channel to exchange requests and acknowledgments for protection switch actions. Even if the protocol is not needed to complete the protection switch actions (as in the case of 1+1 unidirectional operation), it can still be used for other purposes¹¹ and therefore the appropriate codes must be transmitted. The bit assignments for these bytes and the bit-oriented protocol are defined in this section.

5.3.5.1 K1 Byte

This section contains the bit assignments and the APS channel generation rules for the K1 byte carried on the protection line. This byte is used to indicate a request by a channel for a switch action.

5.3.5.1.1 Bit Assignments for the K1 Byte

The K1 byte is divided into two parts. Bits 1 through 4 are used to indicate a request, while bits 5 through 8 are used to indicate which channel is making the request.

Three categories of requests can be indicated using K1 bits 1 through 4:

-
10. It is sufficient to provide provisionable WTR times from 5 to 12 minutes in 1-minute increments.
11. For example, LTE could determine (from the received K1 byte) which of its transmitted OC-N signals is being selected at the far-end LTE, and could provide that information to the user.
-

1. Automatically initiated requests (i.e., SF and SD), which are used to indicate that a failure or degradation condition has been detected on the line associated with the requesting channel
2. An external request (i.e., Lockout of Protection, Forced Switch, Manual Switch, Exercise)
3. A state request (i.e., Wait-to-Restore, Do Not Revert, No Request, Reverse Request), which is used to indicate the state of the APS controller when no other request is indicated.

R5-58 [179] Bits 1 through 4 of the K1 byte shall indicate the current request using the codes listed in Table 5-4.

R5-59 [180] An NE using the 1:n architecture shall provide the capability to provision each working channel and the null channel (for conditions detected on the Protection line) as high or low priority, with low priority as the default. These priorities shall determine which of the listed codes is used for SF and SD requests.

Table 5-4. K1 Byte, Bits 1 through 4: Type of Request

Bit 1234	Automatically Initiated, External, or State Request (Note 1)
1111	Lockout of Protection
1110	Forced Switch
1101	SF - High Priority (Note 2)
1100	SF - Low Priority
1011	SD - High Priority (Note 2)
1010	SD - Low Priority
1001	(not used)
1000	Manual Switch
0111	(not used)
0110	Wait-to-Restore (Note 3)
0101	(not used)
0100	Exercise (Note 4)
0011	(not used)
0010	Reverse Request (Note 5)
0001	Do Not Revert (Note 6)
0000	No Request

Notes:

1. Request priority is in descending order, except that an SF request by the null channel (for an SF condition detected on the protection line) has a higher priority than a Forced Switch (i.e., it is between Lockout of Protection and Forced Switch).
2. High Priority codes apply only to the 1:n architecture.
3. 1+1 LTE provisioned for nonrevertive switching does not transmit Wait-to-Restore.
4. Exercise may not be applicable in some linear APS systems.
5. Reverse Request applies only to bidirectional systems.
6. Only 1+1 LTE provisioned for nonrevertive switching transmits Do Not Revert.

- R5-60** [181] Bits 5 through 8 of the K1 byte shall indicate the number of the channel for which the request is issued, using the codes shown in Table 5-5.

Table 5-5. Channel Number Code Assignments, K1 Bits 5 to 8
(and K2 Bits 1 to 4)

Code	Channel and Notes
0 (i.e., 0000)	Null channel SD and SF requests apply to conditions detected on the protection line. For 1+1 systems, Forced and Manual Switch requests apply to the protection line. Only code 0 is used with the Lockout of Protection request.
1 through 14 (i.e., 0001 through 1110)	Working channels Codes 1 through n apply in a 1:n architecture. Only code 1 applies in a 1+1 architecture. Conditions SD and SF with the provisioned priority (high/low) apply to the corresponding working lines.
15 (i.e., 1111)	Extra traffic channel May exist only when provisioned in a 1:n architecture Only No Request is used with code 15

5.3.5.1.2 K1 Byte Generation

This section contains the K1 byte generation criteria, and the criteria to determine which of the possible state requests to indicate after an automatically initiated request or external request is cleared, or after a WTR state times out.

The K1 byte generation criteria are divided into three parts. These are an evaluation to determine the highest priority local request, a comparison of the highest priority local request with the current local request, and a comparison of the current local request with the remote request. The parts that apply to a particular system depend on the mode of operation.

- R5-61** [182] All local requests [i.e., any locally detected SF or SD conditions, local WTR, Do Not Revert (DNR) or No Request state, or external request from a received switch command] shall be evaluated to determine the highest priority local request based on the order of request priorities in

Table 5-4. If local SF or SD conditions of the same priority have been detected and are still present on different lines at the same time, then the condition with the lowest channel number shall take precedence in the evaluation.

... The highest priority local request shall be compared to the current local request. If the highest priority local request is of higher priority than the current local request (based on the order of priorities in Table 5-4) or if the current local request is no longer a valid request (e.g., the condition or external request that caused it has been cleared) then the highest priority local request shall replace the current local request (i.e., it becomes the new current local request). In all other cases the current local request shall not change.¹²

The LTE's use of the current local request in generating the transmitted K1 byte depends on its mode of operation, as follows.

Bidirectional operation:

R5-62 [183] In the bidirectional mode, the priorities of the current local request and the remote request on the received K1 byte shall be compared according to the order of priorities in Table 5-4. A received Reverse Request shall not be considered in the comparison, because it assumes the priority of the request to which it is responding. The transmitted K1 byte shall be set to indicate a Reverse Request if any of the following are true:

- ... • The remote request is of higher priority
- ... • The requests are of the same priority, they are higher priority than a No Request, and the transmitted K1 byte is already set to indicate Reverse Request
- ... • The requests are of the same priority, they are higher priority than a No Request, the transmitted K1 byte is not set to indicate Reverse Request, and the remote request indicates a lower channel number.
- ... The transmitted K1 byte shall be set to indicate the local request in all other cases.

12. In the bidirectional mode of operation, if the current local request is an SD or SF request that is not being indicated on the K1 byte (i.e., Reverse Request is being indicated instead), then it is acceptable for the highest priority local request to replace the current local request if they have the same priority and the highest priority local request has a lower channel number. In this situation, whether or not the current local request is changed is not critical to the operation of the linear APS system, since the switch requested by the far-end LTE will not be dropped in either case (see R5-62 [183]). What is critical is that if an SD or SF condition on a particular line is still present and is being indicated on the transmitted K1 byte, then it must not be replaced by another SD or SF request with the same priority but a different channel number (i.e., an existing switch that is still needed must not be dropped for another switch with the same priority but a different channel number).

One result of the above requirement is that if LTE is transmitting a particular local request, and then it receives a remote request that has the same priority and the same channel number, it will continue to transmit the local request (and consider the remote request to be a valid response). Another result is that WTR and Do Not Revert (DNR) requests are normally acknowledged by a Reverse Request, while a No Request is acknowledged by a No Request.

Unidirectional operation:

- R5-63** [184] In the unidirectional mode, the transmitted K1 byte shall be set to indicate the current local request (i.e., Reverse Request is never indicated).

When a switch action is no longer requested (i.e., an SD or SF condition clears, an external request is cleared, or a WTR state times out), a new state request is activated and is evaluated according to the rules described above. The new state request depends on the type of switching (i.e., revertive or nonrevertive), the request that is being replaced, and the channel for which that request was issued, as follows.

Working channels at LTE using revertive switching:

- R5-64** [185] For working channels at LTE using revertive switching, when a local condition that caused an automatically initiated switch clears, a local Wait-to-Restore (WTR) state shall be activated.

If the WTR state is the highest priority request, it is indicated on the transmitted K1 byte and maintains the switch of that channel.

- R5-65** [186] A WTR state shall normally time out and become a No Request – null channel (or No Request – Channel 15, if applicable). The WTR timer shall deactivate earlier if the transmitted K1 byte no longer indicates WTR (i.e., when any request of higher priority preempts this state). When the higher priority request is cleared, the preempted WTR state shall not be reactivated. (Note however, that a new WTR state would be required to be activated if the higher priority request was for an automatically initiated switch of a working channel.)

- R5-66** [187] When an external request is cleared, the No Request – null channel (or No Request – Channel 15, if applicable) state shall be activated (i.e., the WTR state is not activated).

The working channel at LTE using nonrevertive switching:

- R5-67** [188] For the working channel at LTE using nonrevertive switching, the selection of the working channel from the protection line shall be maintained by activating a Do Not Revert (DNR) state (instead of a WTR or No Request state). The DNR state shall be deactivated if the transmitted

K1 byte no longer indicates DNR (i.e., when any request of higher priority preempts this state).

Note that if the most recent request was an Exercise request for the working channel, then the selector will be in the released position (see Section 5.3.5.4) and there will be no selection to maintain. Thus, the new request is expected to be No Request – null channel rather than DNR – Channel 1.

Null channel (nonrevertive and revertive):

R5-68 [189] After any request for the null channel is cleared, the No Request – null channel (or No Request – Channel 15, if applicable) state shall be activated.

5.3.5.2 K2 Byte

This section contains the bit assignments and the APS channel generation rules for the K2 byte carried on the protection line. This byte is used to indicate the bridging actions performed at the LTE, and the provisioned architecture and mode of operation.

5.3.5.2.1 Bit Assignments for the K2 Byte

R5-69 [190v2] The bit assignments for the K2 byte shall be as follows (see Section 5.3.5.2.2 for the corresponding K2 byte generation rules):

- ... • Bits 1 through 4 – A channel number using the codes shown in Table 5-5.
- ... • Bit 5 – Indication of architecture (1+1 or 1:n), as provisioned.
- ... • Bits 6 through 8 – Indication of the mode of operation, as provisioned, or non-APS channel uses (i.e., AIS-L, RDI-L).

5.3.5.2.2 K2 Byte Generation Rules

R5-70 [191] For all architectures and modes of operation, bits 1 through 4 of the K2 byte shall be set to indicate:

- ... • The null channel (0) if the received K1 byte indicates the null channel
- ... • The number of the channel bridged onto the protection line in all other cases.

Note that in some situations (e.g., in the bidirectional mode, when the received K1 byte contains a request for a channel that has been locked out by a “Lockout a Working

Channel” command), the null channel may still be the channel bridged onto the protection line even though the received K1 byte does not indicate the null channel.

Also note that although LTE operating in the 1+1 architecture continuously transmits the working channel on the protection line, the working channel is only considered to “bridged” for the purpose of generating the K2 byte if the received K1 byte indicates a request for a valid channel number (i.e., it is not considered bridged for the purpose of generating K2 if the received K1 byte contains any of the invalid channel numbers, 2 through 15).

Similarly, even though LTE operating in the 1:n architecture is allowed to transmit a working channel on the protection line when it is detecting a request for the null channel on the received K1 byte (see Section 5.3.7.1.1), that working channel is also not considered bridged for the purpose of generating the K2 byte. For both the 1+1 and 1:n architectures, non-null channels that are transmitted on the protection line are only considered bridged for the purpose of generating the K2 byte if they are being transmitted in response to a request.¹³ For example, assume 1:n LTE is receiving the No Request – null channel code on K1, and is transmitting working channel 1 on the protection line (and indicating the null channel on K2, as required). If that LTE subsequently receives a request for working channel 2, it should continue to indicate the null channel on the transmitted K2 byte until it completes the requested bridge and changes the code on K2 to ‘0010.’

R5-71 [192] Bit 5 of the K2 byte shall be set to indicate:

- ... • Code ‘0’ if the provisioned (or only supported) architecture is 1+1
- ... • Code ‘1’ if the provisioned (or only supported) architecture is 1:n.

R5-72 [193] Bits 6 through 8 of the K2 byte shall be set to indicate:

- ... • Code ‘101’ if the provisioned mode is bidirectional
- ... • Code ‘100’ if the provisioned (or only supported) mode is unidirectional.

The codes ‘011,’ ‘010,’ ‘001,’ and ‘000’ in K2 bits 6 through 8 are reserved for future use in 1:n drop and insert (nested) protection switching, and the code ‘111’ is used in detecting an incoming AIS-L. Also, **R5-72 [193]** is overridden by the requirements in Section 6.2.1.3.1 when a defect that causes RDI-L (i.e., ‘110’) to be generated is detected on the incoming signal on the protection line.

5.3.5.2.3 Mode of Operation

The K2 byte can be used to determine the protection configuration of the far-end LTE. As discussed in Section 6.2.1.1.6.C, LTE provisioned to operate in any architecture and mode

13. This interpretation is not considered essential to the operation of the linear APS system, and therefore is not listed as a criterion. It is included here for clarification purposes only.

other than the 1+1 unidirectional mode is required to monitor K2 bit 5 and K2 bits 6 through 8 (on the protection line) for APS Mode Mismatch defects and failures.

- R5-73** [194] An APS Mode Mismatch indication resulting from a mismatch of K2 bit 5 shall be used to modify the operation of the 1:1 LTE to interwork with the 1+1 LTE (see Section 5.3.2.3).
- R5-74** [195] An APS Mode Mismatch indication resulting from a mismatch of K2 bits 6 through 8 shall be used by 1+1 LTE provisioned for bidirectional switching to operate unidirectionally, or by 1:n LTE provisioned for unidirectional switching to operate bidirectionally (see Sections 5.3.2.1 and 5.3.2.2).
- R5-75** [196] If LTE stops receiving an indication of the provisioned mode of operation from the far-end LTE, then it shall maintain its current mode of operation.

The near-end LTE will normally stop receiving the far-end LTE's indication of its provisioned mode when the far-end LTE inserts/removes RDI-L on the protection line, or when a regenerator inserts AIS-L on the protection line.

5.3.5.3 Control of the Bridge

The criteria on the bridging of a channel to the protection line depend on the architecture and mode of operation being used, as follows. Also see Figures 5-6 and 5-7.

- R5-76** [197v2] In the 1:n architecture, the channel whose number is indicated on the received K1 byte shall be bridged to the protection line unless the request is invalid (e.g., in the bidirectional mode, the requesting channel is locked out locally).
- R5-77** [198] If a local SF condition is detected on the protection line or a Protection Switching Byte failure is declared, the current bridge shall be maintained if the mode of operation is unidirectional, or shall be released if the mode is bidirectional.
- R5-78** [199] In the 1+1 architecture, the working channel shall be continuously bridged to the protection line.

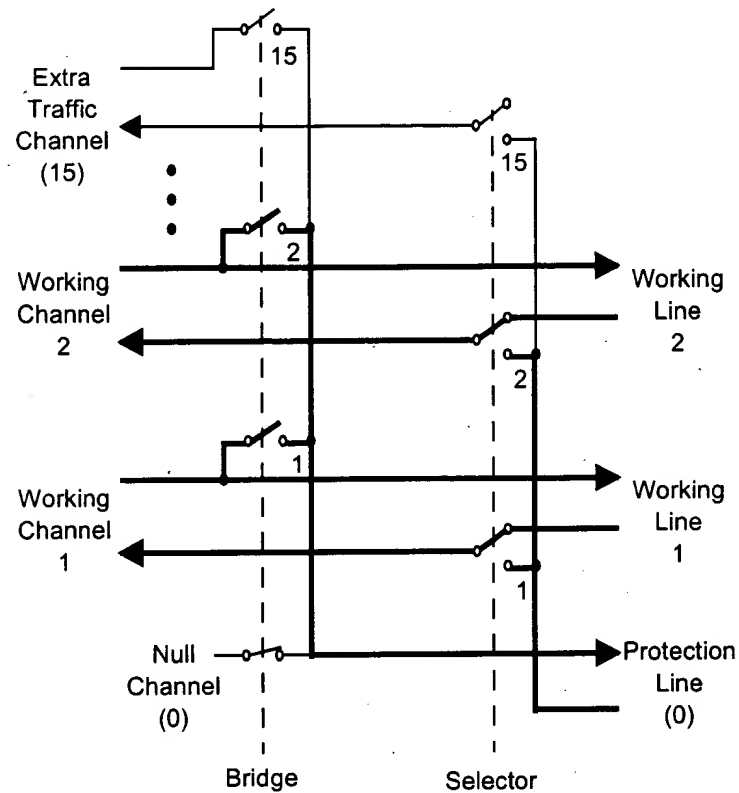


Figure 5-6. Linear APS Switch – 1:n Architecture (in released position)

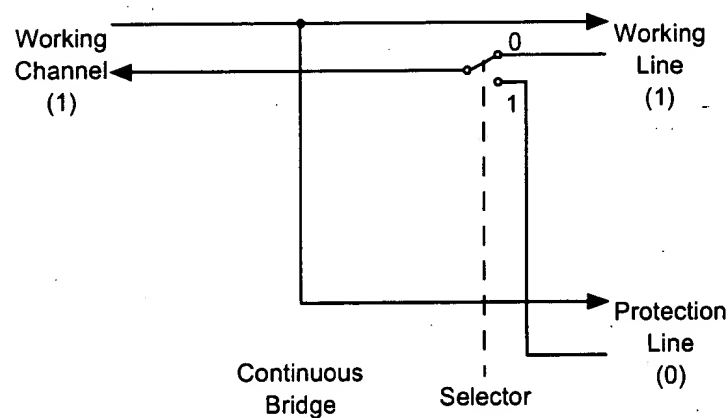


Figure 5-7. Linear APS Switch – 1+1 Architecture (in released position)

5.3.5.4 Control of the Selector

The criteria on the control of the selector depend on the architecture and mode of operation being used. In all architectures and modes except the 1+1 unidirectional mode, the selector is controlled by comparing the channel numbers indicated on the transmitted K1 and received K2 bytes, and the following criteria apply:

- R5-79** [200] In all architectures and modes except the 1+1 unidirectional mode, if there is a match of the transmitted K1 and received K2 bytes, then the indicated channel shall be selected from the protection line, unless one of the following is true (in which case the selector shall be in the released position, see Figure 5-6):
- ... • The match is for the null channel.
 - ... • An Exercise request is indicated on the transmitted K1 byte (unidirectional and bidirectional), or the received and acknowledged K1 byte (bidirectional only).
- R5-80** [201] In the 1:n architecture, the selector shall also be in the released position when there is a mismatch of the channel numbers.¹⁴

In the 1+1 bidirectional mode, the working channel is continuously bridged to the protection line, and therefore the selection of the working channel from the protection line can be made before a match is obtained (i.e., the selector does not have to be in the released position when there is a mismatch). However, if the selection is made before a match is obtained, the match must still be obtained to avoid the declaration of a Channel Mismatch failure (see Section 6.2.1.1.6.B).

In the 1+1 unidirectional mode, the highest priority local request controls the selector, and the following requirement applies:

- R5-81** [202] In the 1+1 unidirectional mode, the working channel shall be selected from the protection line if channel number 1 is indicated on the transmitted K1 byte.

14. The phrase "when there is a mismatch of channel numbers" was intended to mean "after the third of three consecutive received frames containing a channel number (or numbers) in K2 that is different than the transmitted channel number in K1." If an NE meets the Channel Mismatch defect detection objective in Section 6.2.1.1.6.B, then "when there is a mismatch of channel numbers" can be replaced by "when a Channel Mismatch defect is detected". However, if the NE only meets the Channel Mismatch defect requirement, then delaying the release of the selector until the Channel Mismatch defect is detected could result in traffic from one channel being misrouted to a different channel for as long as 50 ms. Therefore, the LTE must release the selector after three frames.

5.3.5.5 Transmission and Acceptance of Bytes K1 and K2

- R5-82** [203] The linear APS protocol shall be carried between LTE in the APS channel (i.e., the K1 and K2 bytes) transmitted on the protection line.

The values transmitted by LTE in the K1 and K2 bytes on the working lines are undefined (except for K2 bits 6, 7, and 8 when RDI-L is being transmitted), and therefore the criteria in Section 3.2 are applicable.¹⁵

- R5-83** [204] A new code on the received K1 and K2 bytes shall replace the current received code if it is received identically in three consecutive frames.

An invalid code is defined as an unused code (see Table 5-4) or a code irrelevant for the specific operation or mode of operation (e.g., a switch request issued for a nonexistent channel).

- R5-84** [205] LTE shall not transmit invalid codes.

- R5-85** [206] If the capability to transport extra traffic on the protection line is provided, the No Request code shall be used to keep the extra traffic channel on the protection line (i.e., the No Request code is the only valid code to transmit with a channel number of '1111').

The preceding requirement implies that the extra traffic channel will be preempted when any request with priority higher than No Request is indicated (i.e., a Lockout of Protection, any request by a working channel to use the protection line, or an SD or SF condition detected on the protection line). One or more of the working channels can be precluded from preempting the extra traffic channel with the "Lockout a Working Channel" control command (see Section 5.3.6.2). However, there is no mechanism defined such that the extra traffic can continue to be transported when an SD condition is detected on the protection line. This issue was discussed at length in GR-253-ILR, Issue ID 253-9. Based on the comments that were received in response to that issue, it was determined that no additional criteria are needed in this area.

It is expected that the LTE at both ends of a linear APS system will normally have consistent views of the priorities (for SD and SF requests) of each of the working and null channels that they jointly support. For example if one LTE in a 1:2 system is provisioned to consider working channel 1 as high priority, and working channel 2 and the null channel as low priority, then the other LTE would normally be expected to be provisioned the same way. In addition, this document and ANSI T1.105.01, *Synchronous Optical Network (SONET) - Automatic Protection Switching*, both indicate that only the low priority SD and

15. It is also acceptable for the K1 and K2 bytes in the signals carried on working lines to carry the APS channel information. However, LTE cannot assume that the APS channel information will be present in the received signal, and therefore it must be capable of ignoring any APS codes received in those bytes.

SF request code are applicable in 1+1 systems. However, the LTE in a 1:n system could easily be provisioned (or misprovisioned) with different views of the channel priorities, and ITU-T Recommendation G.783, *Characteristics of Synchronous Digital Hierarchy (SDH) equipment functional blocks*, indicates that only the high priority codes are applicable in 1+1 systems. To promote compatibility between LTE with different views of the priorities of the various channels, the following objective is applicable:

- O5-86** [917] LTE should not consider an SD or SF request detected on the incoming K1 bits 1 through 4 to be invalid based (solely) on the high/low priority of that request.

For example, LTE that detects a high priority SD request for a channel that it views as low priority should not consider the incoming request to be invalid based solely on the fact that its view of the requesting channel's priority does not match the priority indicated by the incoming request. Note however, that such a request could be considered invalid for other reasons (e.g., in a 1:n system operating in the bidirectional mode, it would be considered invalid if the requesting channel has been locked out using the "Lockout a Working Channel" control command).

In addition, the following criteria are applicable to LTE operating in the 1+1 bidirectional mode.

- R5-87** [207] LTE operating in the 1+1 bidirectional mode and using nonrevertive switching shall consider the WTR code for the working channel to be valid.
- O5-88** [208] LTE operating in the 1+1 bidirectional mode should consider the Do Not Revert code for either the null channel or the working channel as valid.

R5-87 [207] and **O5-88** [208] are needed to facilitate interworking between LTE that uses revertive switching and LTE that uses nonrevertive switching. If the NEs in a 1+1 bidirectional system with one revertive NE and one nonrevertive NE do not conform to **R5-87** [207] and **O5-88** [208], then they will detect and declare unnecessary Protection Switching Byte defects and failures (see Section 6.2.1.1.6.A) after many completed switch requests are cleared. As discussed in **R5-90** [210], the protection line is considered to be in an SF condition when a Protection Switching Byte failure is declared. That SF condition will cause an immediate switch back to the working line, defeating the purpose of both the WTR and DNR states. Conversely, if the NEs conform to the criteria, then no unnecessary failures will be declared and the type of switching used by the system will simply depend on which of the NEs is "in control" (i.e., which NE originated the most recent request that resulted in a completed switch) as described below:

- If the revertive NE is in control, then when the request is cleared that NE will normally send a WTR request (if the original request was an automatically initiated request for a working channel) or a No Request with a channel number of '0' (if it was an external

request or a request for the null channel). The nonrevertive NE will then normally send a Reverse Request (if it receives a WTR) or a No Request (if it receives a No Request). This is exactly the same response as would be expected if the second NE was provisioned for revertive switching rather than nonrevertive switching. Thus, the system uses revertive switching if the revertive NE is in control.

- If the nonrevertive NE is in control, then when the request is cleared that NE will normally send a DNR request (if the original request was for a working channel) or a No Request (if it was for the null channel). The revertive NE will then normally send a Reverse Request (if it receives a DNR) or a No Request (if it receives a No Request). This is exactly the same response as would be expected if the second NE was provisioned for nonrevertive switching rather than revertive switching. Thus, the system uses nonrevertive switching if the nonrevertive NE is in control.

O5-88 [208] is also needed to facilitate interworking between LTE designed to different issues of Bellcore's SONET criteria documents. Some earlier issues of those documents indicated that a nonrevertive NE was supposed to send the DNR code with a channel number of '0' when it was selecting the traffic from the working line and had no higher priority requests. Other issues, including this document, indicate that such an NE is required to send the No Request code (see **R5-68 [189]**).

Receipt of an invalid code or persistently unacceptable codes in the K1 byte results in a Protection Switching Byte defect, and if that defect persists for an extended period, a Protection Switching Byte failure is declared (see Section 6.2.1.1.6.A). Similarly, receipt of an invalid code or persistently unacceptable codes in bits 1 through 4 of the K2 byte results in a Channel Mismatch defect, and if that defect persists for an extended period, a Channel Mismatch failure is declared (see Section 6.2.1.1.6.B).

R5-89 [209] Even when accepted as the current code, an invalid code in K1 shall not result in any immediate protection switching action.

R5-90 [210] For LTE operating in the bidirectional mode, the protection line shall be considered to be in the SF condition when a Protection Switching Byte failure is declared. An SF condition resulting from a Protection Switching Byte failure shall be cleared when the Protection Switching Byte failure is cleared.

For LTE operating in the unidirectional mode, the received request does not affect the transmitted request, so the preceding requirement does not apply.

5.3.6 Linear APS Commands

This section describes the linear APS commands defined to allow the user to perform protection switch actions or to provision the linear APS controller.

5.3.6.1 Switch Commands

A switch command issued at the APS controller interface initiates one external request for evaluation as described in Section 5.3.5.1.2.

- R5-91 [211]** The following switch commands shall be provided, as described.
- ... **Clear** – Clears all of the switch commands listed below, for the channel or channels specified in the command.
 - ... **Lockout of Protection** – Prevents any of the working channels from switching to the protection line by issuing a Lockout of Protection request [unless a request of equal priority (i.e., a Lockout of Protection) is already in effect].
 - ... **Forced Switch of Working (to Protection)** – Switches the specified working channel to the protection line unless a request of equal or higher priority is in effect by issuing a Forced Switch request.
 - ... **Forced Switch of Protection (to Working)** – Switches the working channel back from the protection line to the working line unless a request of equal or higher priority is in effect, by issuing a Forced Switch request for the null channel. This command applies only in the 1+1 architecture.
 - ... **Manual Switch of Working (to Protection)** – Switches the working channel to the protection line unless a request of equal or higher priority is in effect, by issuing a Manual Switch request.
 - ... **Manual Switch of Protection (to Working)** – Switches the working channel back from the protection line to the working line unless a request of equal or higher priority is in effect, by issuing a Manual Switch request for the null channel. This command applies only in the 1+1 architecture.
- R5-92 [212]** When a higher priority local or remote request preempts an external request, the preempted request shall not be retained (i.e., when the higher priority request is cleared, the preempted switch request shall not be reinitiated).

In some situations it could be useful for a previously completed external request to be reinitiated after a higher priority automatically initiated request has preempted it and then cleared (i.e., reinitiate a Manual Switch after an SD or SF request, or a Forced Switch after an SF request for the null channel). However, allowing external requests to be retained and automatically reinitiated would also result in protection switch oscillations in some situations. Therefore **R5-92 [212]** is not expected to be changed in any future issue of this document.

In addition to the required switch commands discussed above, in some applications it may be necessary to support the Exercise command.

CR5-93 [213] LTE capable of operating in the 1+1 bidirectional mode or the 1:n architecture may be required to support the Exercise command.

R5-94 [214] If the Exercise command is supported, it shall cause the LTE to perform as described below.

... **Exercise** – Exercises the protocol for a protection switch of the specified channel, unless a request of equal or higher priority is in effect, by issuing an Exercise request for that channel and checking the response on the APS channel.

As discussed in Section 5.3.5.4, the switch is not actually completed during the Exercise routine (i.e., the selector remains released). Therefore, all actions are the same as those taken for a Manual Switch command, except that the selector remains released.

R5-95 [215] If the Exercise command is supported, it shall be cleared automatically at the end of the Exercise routine or the required switch completion time, whichever is sooner.

Support of the Exercise command is not required or expected for LTE that supports only the 1+1 unidirectional mode. For LTE operating in any of the other possible modes, the following requirement is applicable to facilitate interworking between LTE that supports the Exercise command and LTE that does not support it.

R5-96 [216] LTE that does not support the Exercise command shall consider an incoming Exercise request with a valid channel number to be valid, and shall respond as requested (per Section 5.3.5).

5.3.6.2 Control Commands

Control commands set and modify the linear APS operation. The commands that are currently defined apply only for LTE that supports the 1:n architecture.

R5-97 [217v2] The following control commands shall be supported by 1:n LTE, and shall cause the LTE to perform as described below.

... **Lockout a Working Channel** – Prevents the specified working channel (or channels) from switching to the protection line.

... **Clear Lockout-a-Working-Channel** – Clears the Lockout a Working Channel command for the channel or channels specified in the clear command.

The functionality of the Lockout a Working Channel command depends on the switching mode (i.e., unidirectional or bidirectional) being used. If the switching mode is bidirectional, then the lockout is also bidirectional. Specifically, local requests for the

locked out channel are not considered in the K1 byte generation process, and remote requests for that channel are not acknowledged (i.e., they also are not considered in the K1 byte generation process, and the requested bridge is not performed or indicated in K2 bits 1 through 4). Similarly, if the switching mode is unidirectional, then the lockout is also unidirectional. That is, local requests for the locked out channel are not considered in the K1 byte generation process, but remote requests for that channel are acknowledged (i.e., the requested bridge is performed and indicated in K2 bits 1 through 4).

Note that in the bidirectional mode, a Lockout a Working Channel command for a particular channel must be applied at both ends of the SONET Line for proper operation.

5.3.7 Switch Operation

This section describes the operation of the linear APS system and provides an example of the operation in a tabular format.

5.3.7.1 1:n Architecture

5.3.7.1.1 1:n Bidirectional Mode

Table 5-6 gives an example of the protection switching actions between two multiplexer sites (denoted by A and C) of the 1:n bidirectional linear APS system shown in Figure 5-8. A description of the operation follows the table.

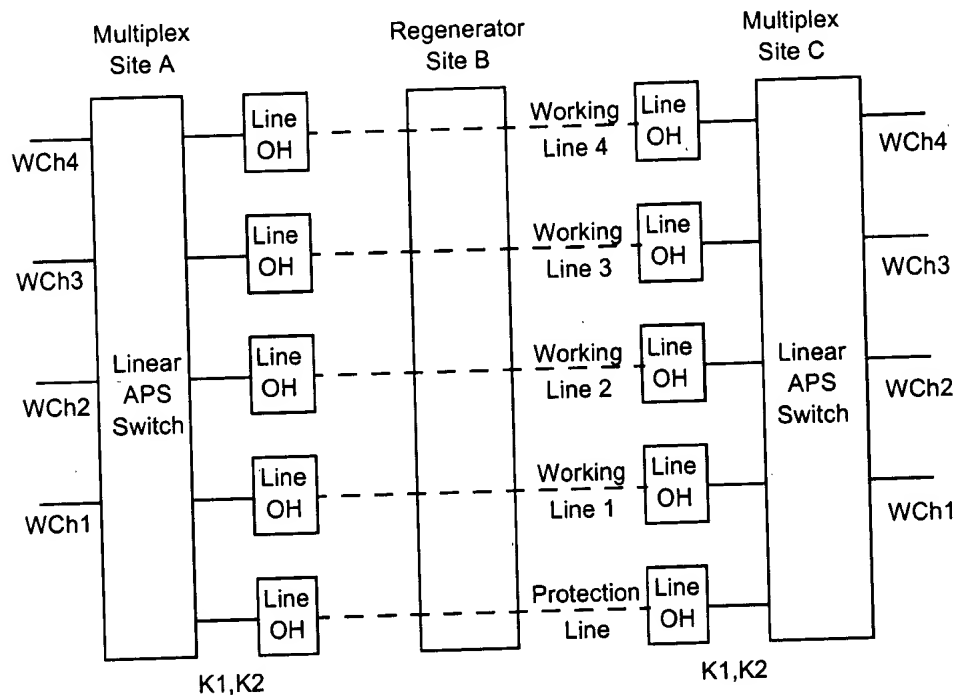


Figure 5-8. 1:n Linear APS Architecture Example

Table 5-6. 1:n Bidirectional Switching Example

Failure Condition or Controller State	APS Bytes				Action	
	C → A		A → C		At Site C	At Site A
	K1 Byte	K2 Byte	K1 Byte	K2 Byte		
No Failures (Protection Line is not in use)	00000000	00001101	00000000	00001101	WCh 3 is transmitted on the protection line to provide a valid signal. Selector is released.	WCh 4 is transmitted on the protection line to provide a valid signal. Selector is released.
Working Line 2 Degraded in Direction A → C	10100010	00001101	00000000	00001101	Failure detected. Request WCh 2 bridge - SD.	
	10100010	00001101	00100010	00101101		Bridge WCh 2. Send Reverse Request for WCh 2 bridge.
	10100010	00101101	00100010	00101101	Select WCh 2. Bridge WCh 2.	
	10100010	00101101	00100010	00101101		Select WCh2. Bidirectional switch completed.
Working Line 1 Failed in Direction C → A (This preempts the WCh 2 Switch)	10100010	00101101	11000001	00101101		Failure detected. Request WCh 1 bridge - SF. Release WCh 2 selection.
	00100001	00011101	11000001	00101101	Bridge WCh 1. Send Reverse Request for WCh 1 bridge. Release WCh 2 selection.	
	00100001	00011101	11000001	00011101		Select WCh 1. Bridge WCh 1.
	00100001	00011101	11000001	00011101	Select WCh 1. Bidirectional switch completed.	

Table 5-6. 1:n Bidirectional Switching Example (Continued)

Failure Condition or Controller State	APS Bytes				Action	
	C → A		A → C		At Site C	At Site A
	K1 Byte	K2 Byte	K1 Byte	K2 Byte		
Working Line 1 Repaired (Working Line 2 still Degraded)	00100001	00011101	01100001	00011101		Send Wait-to-Restore for WCh 1.
	10100010	00011101	01100001	00011101	Request WCh 2 bridge - SD. Release WCh 1 selection.	
	10100010	00011101	00100010	00101101		Bridge WCh 2. Send Reverse Request for WCh 2 bridge. Release WCh 1 selection.
	10100010	00101101	00100010	00101101	Bridge WCh 2. Select WCh 2.	
	10100010	00101101	00100010	00101101		Select WCh 2. Bidirectional switch completed.
Working Line 2 Repaired	01100010	00101101	00100010	00101101	Send Wait-to-Restore for WCh 2.	
Wait to Restore Expired (No Failures)	00000000	00101101	00100010	00101101	Drop WCh 2 request. Release WCh 2 selection.	
	00000000	00101101	00000000	00001101		Release WCh 2 bridge. (Transmit WCh 4 on the protection line.) Drop WCh 2 request. Release WCh 2 selection.
	00000000	00001101	00000000	00001101	Release WCh 2 bridge. (Transmit WCh 3 on the protection line.)	
No Failures (Protection Line is not in use)	00000000	00001101	00000000	00001101	WCh 3 is transmitted on the protection line to provide a valid signal. Selector is released.	WCh 4 is transmitted on the protection line to provide a valid signal. Selector is released.

When the protection line is not in use, the null channel is indicated on the K1 bytes. The null channel is also indicated on the K2 bytes, according to K2 byte generation rules. However, any working channel may actually be transmitted on the protection line by the head end. The tail end does not assume or require any specific channel. In the example in

Table 5-6, working channel (WCh) 3 is transmitted at site C, and WCh 4 is transmitted at site A.

When an SF or SD condition is detected or a switch command is received at the tail end of a line, the APS controller compares the priority of this new condition with the request priority of the channel (if any) on the protection line. The comparison includes the priority of any remote request on the received K1 byte. If the new request is of higher priority, the K1 byte is loaded with the request and the number of the channel asking for use of the protection line. In the example, an SD condition is detected at C on working line 2, and an SD request is sent on the K1 byte to A.

At the head end, when this new incoming K1 byte has been verified (i.e., received identically for three consecutive frames) and evaluated by the APS controller, the outgoing K1 byte is sent with a Reverse Request and the appropriate channel number as a confirmation for that channel to use the protection line, and to request a bridge at the tail end for the same channel. A Reverse Request is returned for all requests except a No Request. This clearly identifies which end originated the switch request. If the head end also originates an identical request (not yet confirmed by a Reverse Request) for the same channel, then both ends continue transmitting the identical K1 byte and perform the requested switch action.

Also, at the head end, the indicated channel is bridged to the protection line. When the channel is bridged, the K2 byte is set to indicate the number of the channel on the protection line.

At the tail end, when the channel number on the received K2 byte matches the number of the channel requesting the switch, that channel is selected from the protection line. This completes the switch to the protection line for one direction. The tail end also performs the bridge as requested by the K1 byte and indicates the bridged channel on the K2 byte. The head end completes the bidirectional switch by selecting the channel from the protection line when it receives a matching K2 byte.

Normally, a switch is completed in 50 ms. If the switch cannot be completed because an appropriate code is not returned in the K1 byte sent by the head end to the tail end, or because one end does not perform and indicate the appropriate bridge, then a Protection Switching Byte defect and/or a Channel Mismatch defect is detected. If a one of these defects persists for 2.5 seconds, then the appropriate failure is declared. In the case of a Protection Switching Byte failure, the end declaring the failure changes its transmitted K1 byte to indicate an SF on the protection line.

Table 5-6 further illustrates how an existing switch is preempted by a higher priority request. In the example, an SF condition on WCh 1 preempts the WCh 2 switch. The selectors are temporarily released before selecting WCh 1, because of the temporary channel number mismatches on the transmitted K1 and received K2 bytes. The example also illustrates the switch back to WCh 2 after the failure on working line 1 is repaired.

When the switch is no longer needed (e.g., the failure on working line 2 is repaired) and the WTR time has expired, the tail end indicates a No Request – null channel on the K1 byte. This releases the selector because of a channel number mismatch. The head end then releases the bridge and replies with the same K1 byte and the null channel number on the K2 byte. The selector at the head end is also released because of a channel number mismatch. Receiving the number for the null channel on the K1 byte causes the tail end to release the bridge. Because the K2 bytes now indicate the null channel, and that matches the null channel numbers on the K1 bytes, the selectors remain released, any Channel Mismatch defects are terminated, and restoral is completed.

5.3.7.1.2 *1:n Unidirectional Mode*

All actions are as described in Section 5.3.7.1.1 for the bidirectional mode, except that the unidirectional switch is completed when the tail end selects the channel for which it issued a request (on the K1 byte) from the protection line. This difference in operation is obtained by not considering remote requests in the priority logic and, therefore, by not issuing Reverse Requests.

5.3.7.2 *1+1 Architecture*

5.3.7.2.1 *1+1 Bidirectional Mode*

For 1+1 bidirectional systems, the K1 and K2 bytes are exchanged as Section 5.3.7.1.1 describes to complete a switch, except that the head end maintains a continuous bridge of the working channel to the protection line, and therefore separate bridging actions are not performed for each request.

In addition, each end is allowed to switch immediately, before receiving a bridge confirmation from the other end. However, the channel matches must still be obtained to avoid Channel Mismatch failures.

Finally, for revertive switching the restoral takes place as described in Section 5.3.7.1.1. For nonrevertive switching (assuming the working channel is being selected from the protection line) when the working line is repaired or a switch command is cleared, the tail end maintains the selection and indicates Do Not Revert for WCh 1. The head end also maintains the selection and continues indicating Reverse Request. The Do Not Revert is removed only when preempted by a failure condition or an external request.

5.3.7.2.2 1+1 Unidirectional Mode

For 1+1 unidirectional switching, the working channel is continuously bridged to the protection line, and the channel selection is based on only the local conditions and requests. Therefore, each end operates independently of the other end, and the K1 and K2 bytes are not needed to coordinate switch actions. However, the K1 byte is still used to inform the other end of the local action, and the K2 byte is set to indicate that the K1 byte is being received (i.e., by indicating the same channel number as the received K1), and to inform the other end of the provisioned architecture and mode of operation.

5.4 Network Synchronization

SONET uses the existing synchronization network as described in GR-436-CORE and ANSI T1.101. This section discusses applications with respect to SONET NEs and the service provider synchronization networks, timing modes for SONET NEs, criteria for the internal SONET clocks, criteria for SONET-based timing distribution, timing reference switching criteria, and criteria for the use of synchronization status messages. This section also indicates which of the criteria in GR-1244-CORE are applicable to SONET NEs.

Note that in this section “OC-N” and “STS-N electrical” are generally used to identify the high-speed SONET “interfaces”¹⁶ to a particular NE, and “OC-M” and “STS-M electrical” are used to identify any low-speed or tributary SONET “interfaces”. For example, in Figure 5-9 the ADM’s OC-12 “interfaces” are OC-N “interfaces” and its OC-3 “interfaces” are OC-M “interfaces”.

5.4.1 SONET NE Clock Applications

The existing service provider synchronization networks are expected to evolve to an OC-N based network for the transport of timing references between locations. The reliability and overall accuracy of transported timing references will be enhanced by the use of OC-N based synchronization networks. Conversely, planning the interoffice synchronization network and avoiding timing loops is much more challenging with the deployment of SONET. Furthermore, DS1s carried on SONET are not recommended for network

16. Throughout Section 5.4, an OC-N (or OC-M) “interface” is defined to include all of the OC-N (or OC-M) signals that are part of a “line APS” system (if supported). In turn, line APS is used to refer to protection switching architectures where an NE receives two or more OC-N (or OC-M) signals from a peer NE, and one of those signals is used to provide SONET Line layer protection for the working channel(s) carried on the other signal(s). This includes systems that support linear APS (see Section 5.3), and span-switching in 4-fiber bidirectional line switched rings (see GR-1230-CORE). If no line APS is supported for a particular pair of OC-N or OC-M signals (e.g., for the OC-N signals that connect two adjacent nodes in a UPSR), then the “interface” includes only those two (i.e., one incoming and one outgoing) signals. In this section, quotation marks will be used to distinguish between an “interface” as defined here and other types of interfaces that are commonly defined (e.g., a point in the transmission medium at which the physical layer characteristics of a signal are specified and measured).

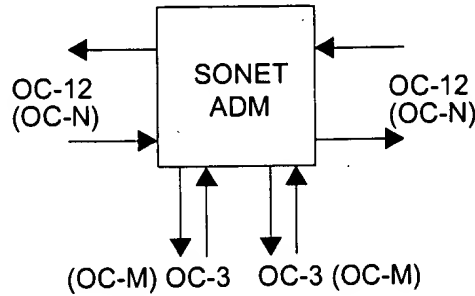


Figure 5-9. OC-N and OC-M Example

synchronization distribution to BITS clocks, because these signals will not meet ANSI T1.101 synchronization interface specifications.

SONET NEs are required to have internal clocks of ± 20 -ppm minimum free-run accuracy. Based on the work done in T1X1 on ANSI T1.105.09, *SONET: Network Element Timing and Synchronization*, clocks that are contained in NEs that support SONET Line terminating functions, and that meet only this minimum accuracy requirement are called SONET Minimum Clocks (SMCs).

Many synchronization-related criteria are application-specific, particularly for SONET ADMs. External or loop-timing may be appropriate for a TM,¹⁷ while external, line, or through-timing may be appropriate for an ADM. Protection switching schemes and dropped OC-M and STS-M electrical signals further complicate the selection of synchronization options. In general, most of the information and criteria applicable to a SONET ADM appear in this document; however, GR-496-CORE also contains several criteria related to the timing options that should be supported. Also note that other NEs, such as DCSs and SONET regenerators, have specific synchronization criteria that are covered in the NE-specific GRs, TRs, and TAs.

The following general statements describe conditions necessary for synchronization and SONET networks to be compatible:

- Where BITS timing is available, SONET NEs are externally timed from the BITS clock.
- Where no BITS timing is available, SONET NEs are timed from a received OC-N (or OC-M) signal.
- External-timing references to a SONET NE are from a BITS clock of stratum 3 or better quality.

17. In this section, ADMs in the terminal configuration are referred to as Terminal Multiplexers (TMs), while ADMs in the add-drop configuration are simply referred to as ADMs.

- Timing signals delivered to the synchronization network from a SONET NE are derived directly from a terminating OC-N (or OC-M).

5.4.1.1 Physical Interface to Synchronization Network

The physical interface for synchronization signals is important so that SONET NEs can be easily integrated into the BITS plan. BITS clocks provide two types of timing outputs: DS1 and Composite Clock (both of which are balanced signals). As described in Section 5.4.3.1, DS1 signals are the usual timing reference signals for most SONET NEs, but Composite Clock (CC) is required for SONET NEs with DS0 interconnections. In addition, the derived DS1 signals from SONET NEs may be used as references for the BITS clock.

Section 3.2.2 of GR-1244-CORE describes various criteria for the physical interface to the synchronization network, such as termination requirements and wire wrap terminals.

R5-98 [918] The physical interface between the SONET NE and the synchronization network shall meet the criteria in Section 3.2.2 of GR-1244-CORE.

5.4.1.2 TDEV and MTIE Measurements

The discussion of the TDEV and MTIE parameters that appeared in this section in Issue 1 of this document now appears in GR-1244-CORE.

5.4.2 Synchronization Status Messages

Synchronization status messages have been defined as a nibble (bits 5 to 8) in the S1 byte of the SONET line overhead and as a bit-oriented message in the Extended Superframe Format (ESF) data link of DS1 signals. These messages contain clock quality information that allows SONET NEs to select the most suitable synchronization reference from the set of available references. The purpose of these messages is to allow SONET NEs to reconfigure their synchronization references autonomously while avoiding the creation of timing loops. **However, it is critical to realize that the use of synchronization status messages alone will not preclude the creation of timing loops. Synchronization engineering following the guidelines in GR-436-CORE is still required.**

Table 5-7 lists the synchronization status messages that have been defined at this time for the S1 byte and the ESF DS1 format. This table is based on T1 Technical Report #33, *A Technical Report on Synchronization Network Management Using Synchronization Status Messages*. That report is currently under review in T1X1.3 in a number of areas (e.g.,

possible definition of additional messages, the use of the DUS message), and this document may be changed or expanded in response to that work.

Table 5-7. Synchronization Status Message Definitions

Description	Acronym	Quality Level	DS1 ESF Data Link Code Word ^a	S1 bits 5678 ^b
Stratum 1 Traceable	PRS	1	00000100 11111111	0001
Synchronized - Traceability Unknown	STU	2	00001000 11111111	0000
Stratum 2 Traceable	ST2	3	00001100 11111111	0111
Transit Node Clock Traceable ^{c, d}	TNC	4	01111000 11111111	0100
Stratum 3E Traceable ^c	ST3E	5	01111100 11111111	1101
Stratum 3 Traceable	ST3	6	00010000 11111111	1010
SONET Minimum Clock Traceable	SMC	7	00100010 11111111	1100
Stratum 4 Traceable	ST4	8	00101000 11111111	N/A
DON'T USE for Synchronization	DUS	9	00110000 11111111	1111
Reserved for Network Synchronization Use	RES	user assignable	01000000 11111111	1110

Notes:

- The ESF synchronization status messages are transmitted rightmost bit first.
- The S1 bits are transmitted leftmost bit first.
- The codes listed for TNC and ST3E clocks (as well as the inclusion of those types of clocks in this table) are based on the information contained in the drafts of ANSI T1.101 and T1.403 that were available when this document was issued. These codes are currently not expected to be changed during the standards approval process; however, the user should be aware that that possibility exists.
- Except for the purpose of defining synchronization status messages, TNC clocks are not considered in this document.

Synchronization status messages in the S1 byte can provide the following benefits:

- automatic reconfiguration of line-timed rings
- improved reliability of interoffice timing distribution
- trouble-shooting of synchronization-related problems.

Therefore, support for synchronization status messages is required for (almost) all line-side SONET signals.

- R5-99** [224v2] SONET LTE shall generate and provide the capability to process the synchronization status messages listed in Table 5-7 on bits 5 through 8 of the S1 bytes of all signals at SONET line-side "interfaces" (except for lines 2 through n at an OC-N "interface" where 1:n linear APS is being used).

In addition, synchronization status messages are useful as a troubleshooting tool for drop-side signals at network interfaces (such as at a customer premise location), even if those signals are not required to be used as timing references.

- R5-100** [225v3] SONET LTE shall generate the synchronization status messages listed in Table 5-7 on bits 5 through 8 of the S1 bytes of all signals at SONET drop-side "interfaces".

- CR5-101** [1048] SONET LTE may be required to provide the capability to process the synchronization status messages listed in Table 5-7 on bits 5 through 8 of the S1 bytes of all signals at SONET drop-side "interfaces".

In the above criteria, the use of the words "generate" and "provide the capability to process" is intended to indicate that synchronization status messages must always be generated on outgoing signals, but may not need to be processed on incoming signals. This is similar to the criteria for various other overhead bits and bytes, such as those used for REI-L, REI-P and REI-V (e.g., the M1 byte). Those bits and bytes are required to always be generated on originating line and path signals, but are only required to be processed if a far-end PM accumulation feature is active for the incoming line or path. [Note however that in the case of synchronization status messages, a user can effectively override the required generation of those messages at any desired "interfaces" by provisioning them to be set to the DUS message (see **R5-209** [324]).]

As for the processing of the S1 byte contained in incoming SONET signals, the following criteria apply.

- CR5-102** [1049] A SONET NE that contains LTE and supports line-timing (or through-timing), or that is capable of being provisioned to derive DS1s from the incoming signals at one or more of its SONET "interfaces", may be required to be able to be provisioned by the user to ignore the incoming S1 byte at its provisioned reference and/or derived DS1 source "interfaces".

- R5-103** [1050] If a SONET NE allows the user to disable the processing of the S1 byte at its SONET "interfaces" that are provisioned as timing references or derived DS1 sources, the default shall be that the S1 byte is processed.

In general it is assumed that in any particular application an NE will either use the messages received on all of its provisioned reference and derived DS1 source "interfaces", or not use the messages on any of them. Therefore it would be sufficient to provide the capability to disable the use of the incoming synchronization status messages on a per-NE basis. However, it is also acceptable for that capability to be provided on a per-reference basis, or a per-derived DS1 or per-derived DS1 source basis. In addition, note that if the use of the incoming messages has been disabled, then various other criteria (or parts of criteria) in this document will not be applicable (e.g., **RS-192 [302v2]** on considering a reference failed or unavailable based on its synchronization status message). Finally, note that for interfaces that are not or cannot be provisioned as timing references or derived DS1 sources, the only defined use of the incoming synchronization status messages is that the NE should allow the user to retrieve those messages. That capability is supposed to be provided independent of whether processing has been enabled or disabled (see **OS-205 [320v3]**), and therefore it is not necessary for an NE to support the capability to disable processing at those interfaces.

It is important to note that interoperability issues exist in environments when S1 synchronization status messaging is not uniformly supported (e.g., because of an embedded base that was deployed before that capability was developed). Specifically, the creation of timing loops cannot be precluded in certain scenarios. For example, in a line-timed ring where some NEs support synchronization messaging and others do not, a timing loop may be created if the NEs are allowed to reconfigure their timing sources autonomously. This is mainly due to the fact that equipment that does not support S1 messages will generate an all-zeros code in the S1 nibble, which corresponds to the "Synchronized - Traceability Unknown" message. Careful engineering is required in mixed environments to ensure that timing loops are not created.

Messages in the DS1 ESF data link are not needed to realize all of the benefits listed above. Therefore, a SONET NE is only conditionally required to support DS1 ESF messages (see **CR5-171 [287v2]**). However, some service providers believe that the use of DS1 ESF synchronization status messages will add robustness to their interoffice synchronization distribution networks. Therefore, support for DS1 ESF synchronization status messages will be required by some service providers. **It is important to note that the benefit of preventing timing loops for interoffice synchronization distribution (particularly for SONET rings) can only be realized if synchronization status messaging is supported by all TSGs and SONET NEs. Additionally, even if synchronization status messaging is implemented for all SONET signals and ESF DS1s, the creation of timing loops could still occur in certain configurations or situations.** Therefore, careful synchronization planning is still necessary.

T1 Technical Report #33 states that the ESF messages should be sent at least once every 15 minutes, but that they may be sent continuously when the DS1 is a non-traffic carrying, synchronization-only signal. In all the applications for the ESF messages described in this document (i.e., external-timing references from a BITS clock and the derived DS1 outputs from a SONET NE) the DS1 is dedicated to synchronization, so the ESF messages are sent continuously (see Section 5.4.5.2.2).

Network providers may choose a “simplified” synchronization message set with no DS1 ESF synchronization status messages. In that case, the full DS1 ESF message set is replaced with only two possibilities: either a framed DS1 signal or a DS1 AIS signal (i.e., unframed all-ones). The full benefits of synchronization messaging in line-timed rings can be realized when the “simplified” message set is implemented.

The S1 message for “Synchronized – Traceability Unknown” was intentionally selected to be an all-zeros pattern because that is the preferred way to populate a field that is unused (see Section 3.2). Therefore, the expected output of the NE when synchronization status messaging is not supported matches the message that indicates “Synchronized – Traceability Unknown.” The “Synchronized – Traceability Unknown” definition for the ESF data link could not be assigned using the same logic. The “Synchronized -Traceability Unknown” message will be prevalent throughout the network when synchronization status messaging is first implemented. If the embedded base of SONET NEs and BITS clocks are upgraded to support synchronization messaging, the prevalence of the “Synchronized – Traceability Unknown” message will diminish.

Bellcore has not assigned a use for the “Reserved for Network Synchronization Use” message.

R5-104 [222] The “Reserved for Network Synchronization” message shall be treated as a “DON’T USE for Synchronization” message at intercarrier interfaces.

R5-105 [223] Any proprietary use of the “Reserved for Network Synchronization Use” message shall be clearly documented as per Section 3.2.

Synchronization status messages impact criteria in many subsections of Section 5.4. Specifically, external-timing mode criteria, derived DS1 criteria, and reference switching criteria are affected. Also, a section specific to synchronization status message validation and generation is included.

5.4.3 SONET Timing Modes

This section describes the four timing modes for SONET NEs: external, line, loop, and through. The timing mode determines the source of timing for OC-N/OC-M, STS-N/STS-M electrical, and synchronous DS1 signals transmitted by the NE, along with the STS and VT SPEs created by that NE (with the possible exception of VT1.5 SPEs containing byte-synchronously mapped DS1s, as discussed in Section 3.4.1.1). Specific documents that describe the application of an NE identify which timing mode or modes the NE is required to support.

R5-106 [226] For NEs that support the external-timing mode, that mode shall be the default timing mode.

- R5-107** [227] For NEs that support automatic switching between timing modes, the switching shall be revertive.

Note that the above switching requirement applies to *timing mode* switching and not to *reference* switching or *hardware* switching. In general, NEs are provisioned to use only a single timing mode and this requirement would not be pertinent. For example, it is *not* expected that an NE would be provisioned to enter line-timing if all of its external references fail. As discussed in GR-436-CORE, maintaining and administering a network (especially avoiding the creation of timing loops) becomes difficult if this type of switching is allowed. The one case where timing mode switching does apply is for NEs that switch from through-timing to line-timing when an OC-N line fails, as per Section 5.4.3.4.

Sections 5.4.3.2, 5.4.3.3, and 5.4.3.4 describe methods of recovering timing from a terminating OC-N. Use of one of these timing modes is necessary when no external-timing reference is available. The ability to recover timing from terminating STS-N electrical signals is not required.

5.4.3.1 External Timing

In accordance with the BITS concept, external timing from a BITS clock is the preferred mode of synchronizing SONET NEs that contain LTE. Figure 5-10 illustrates external timing.

- R5-108** [919] SONET NEs with external-timing interfaces shall meet the criteria in Section 3.2.1 of GR-1244-CORE.

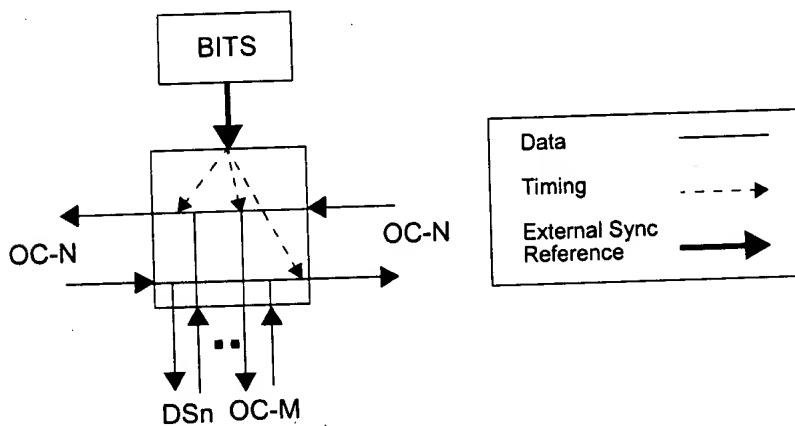


Figure 5-10. External-Timing Mode Example

5.4.3.2 Line-Timing

Figure 5-11 illustrates the line-timing mode for an ADM. Some suppliers have suggested that an alternative to line-timing would be to use a derived DS1 (defined in Section 5.4.5.1) as an external reference. This may be acceptable in some applications, but cannot be called line-timing.

R5-109 [920] SONET NEs with line-timing “interfaces” shall meet the criteria in Section 3.2.3 of GR-1244-CORE.

Note that requirement **R1244-16** allows (for example) an ADM to switch between its “east” and “west” “interfaces” if the “interface” currently used as the timing reference becomes unavailable as defined in Section 5.4.6. However, provisioning multiple “interfaces” as synchronization references must be done with care to avoid the creation of timing loops. Refer to GR-436-CORE for more information about synchronization planning and administration. In general, NEs without S1 synchronization status messaging capabilities should have only one OC-N “interface” provisioned as a reference.

While it is not required (or recommended by Bellcore) that an NE that supports line APS also support the provisioning of working line 1 and the protection line at a single SONET “interface” as separate references, such provisioning is allowed and may be supported by some NEs. In addition, some of those NEs may not allow the user to provision a line APS “interface” as a single reference. In that case, the following requirement applies:

R5-110 [1051] If an NE that supports line APS does not support provisioning of an OC-N/M “interface” as a single reference, that fact shall be clearly documented.

Note that even if an NE supports the provisioning of working line 1 and the protection line as separate references (or as separate derived DS1 sources), it still must meet the applicable criteria related to the generation of the DUS synchronization status message (e.g., **R5-219 [330]**, **R5-211 [326v2]**).

CR5-111 [239] In some applications the NE may be required to provide the user the capability to provision a line-side OC-M “interface” as a synchronization source.

A typical application where this would apply is one where the timing distribution path is through a low-speed or tributary “interface” on an ADM. For example, a ring on a customer’s campus environment that has a low-speed SONET connection to a central office (to carry the traffic going on and off of the campus) may need to receive timing from the central office via that low-speed SONET “interface”.

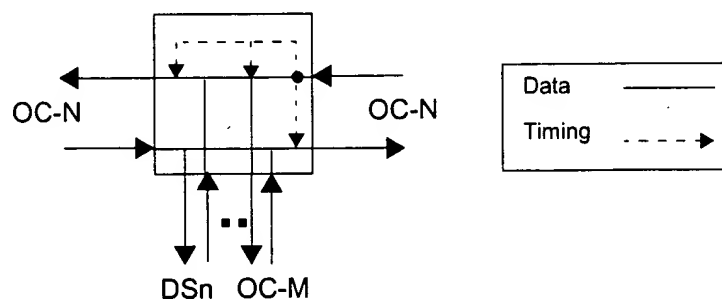


Figure 5-11. Line-Timing Mode Example

5.4.3.3 Loop-Timing

Loop-timing is a special case of line-timing. It applies to NEs that have only one OC-N “interface” (see Figure 5-12).

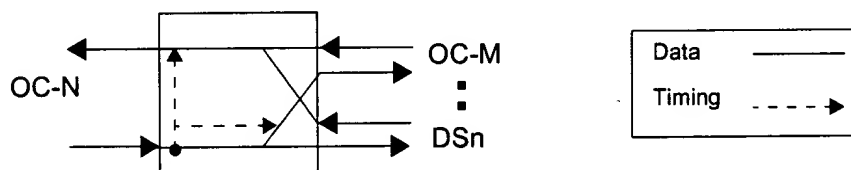


Figure 5-12. Loop-Timing Mode

5.4.3.4 Through-Timing

Through-timing is not a recommended timing mode for SONET NEs that contain LTE (e.g., ADMs). For such NEs, the through-timing mode can be a very complex timing scheme. For example, the user may find it unclear what the timing source is for a low-speed “interface”, or for the protection line at an NE that supports line APS (particularly if a 1:n APS architecture is being used). However, through-timing is the required timing mode for regenerators (see TR-NWT-000917), and is supported by some other types of deployed SONET NEs. Therefore, aspects of the timing mode must be specified.

Figure 5-13 illustrates the through-timing mode for an ADM. The following criteria apply to NEs configured for through-timing.

- R5-112** [240v2] When an NE is through-timed, the transmitted signals at the “west” OC-N or STS-N electrical “interface” shall be timed by the terminating signals at the “east” OC-N or STS-N electrical “interface”, and the transmitted signals at the “east” “interface” shall be timed by the terminating signals at the “west” “interface”.
- R5-113** [241] An ADM that supports through-timing shall provide the user with the capability to provision for automatic switching to line-timing using the non-failed OC-N “interface” when one of the OC-N “interfaces” becomes unavailable as a reference, as defined in Section 5.4.6.
- R5-114** [242] If a through-timed ADM has OC-M, STS-M electrical, or synchronous DS1 interfaces, the timing source for these outputs, as a group or individually, shall be user-provisionable from either of the OC-N “interfaces”.

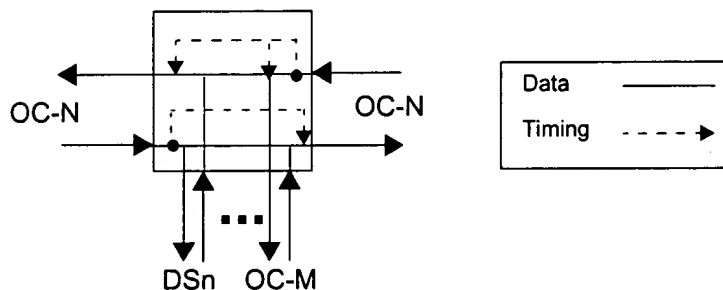


Figure 5-13. Through-Timing Mode Example

5.4.4 SONET Internal Clock

This section provides internal clock criteria. All clock performance characteristics are measured at the synchronous outputs (e.g., OC-N outputs) of the NE.

Note that the criteria in this section are intended for SONET NEs that have stratified clocks or SMCs, and that some SONET NEs are not expected to provide such clocks. In particular, SONET NEs that do not contain LTE (e.g., regenerators) need only provide a free-running clock with ± 20 ppm accuracy (see TR-NWT-000917). That clock is used to generate AIS-L when the NE detects an incoming signal defect (i.e., an LOS or LOF), and therefore is not receiving a valid timing reference.¹⁸ For such NEs, the Category II jitter transfer

18. Although the accuracy requirements for an SMC and a regenerator's clock are identical (± 20 ppm), the regenerator's clock is not required to meet the other performance criteria applicable to an SMC and therefore it is not considered to be an SMC.

requirement in Section 5.6.2.1.2 is applicable, rather than the wander transfer and generation criteria in this section.

5.4.4.1 Stratum Clocks for SONET Applications

Stratum clock criteria are contained in GR-1244-CORE. Some applications require that SONET NEs provide clocks that meet, as a minimum, stratum 3, 3E, or 2 criteria.

- CR5-115 [243v2]** Some service providers may require stratum 3 clocks for SONET NEs used in applications that do not explicitly require stratum 3 clocks. |
- CR5-116 [244v2]** Some providers may require NEs to provide stratum 3E clocks in certain applications, such as NEs that serve as timing distribution hubs. |
- CR5-117 [245v2]** Some providers may require NEs to provide stratum 2 clocks in certain applications, such as NEs that serve as timing distribution hubs. |
- R5-118 [246]** A stratum 3, 3E or 2 clock in a SONET NE shall meet the following criteria in GR-1244-CORE:
- ... • minimum free-run accuracy (GR-1244-CORE, Section 5.1)
 - ... • holdover stability (GR-1244-CORE, Section 5.2)
 - ... • pull-in/hold-in range (GR-1244-CORE, Section 3.5)
- R5-119 [247]** A stratum 3E or 2 clock in a SONET NE shall meet the following requirements in GR-1244-CORE:
- ... • wander tolerance (GR-1244-CORE, Section 4.3)
 - ... • wander transfer (GR-1244-CORE, Section 5.4)
- R5-120 [921]** A stratum 3 clock in a SONET NE shall meet the SMC wander transfer requirements in Section 5.4.4.2.4 (of this document).

Note that the TDEV specifications for the input signals in the SMC wander transfer requirements in Section 5.4.4.2.4 define the minimum wander tolerance for stratum 3 clocks in SONET NEs.

5.4.4.2 SONET Minimum Clock Applications

In general, the following criteria apply to the clocks in all SONET NEs that contain LTE and do not contain a stratum clock (i.e., to SMCs as defined in T1.105.09). Note however, that Section 5.4.4.2.4 also applies to stratum 3 clocks, as per **R5-120 [921]**.

5.4.4.2.1 *Free-Run Accuracy for SMCs*

Free-run accuracy is defined in GR-1244-CORE.

- R5-121** [248] The minimum free-run accuracy of an SMC shall be ± 20 ppm.

5.4.4.2.2 *Holdover for SMCs*

- R5-122** [249v3] A SONET NE containing an SMC shall be capable of entering holdover when all of its timing references are determined to be failed (as per Section 5.4.6) or contain the "DON'T USE for Synchronization" synchronization status message.

Holdover is defined in GR-1244-CORE. Note that for a failure resulting from a frequency offset of a reference, the NE may enter free-run if the holdover value has been corrupted.

In general, a phase transient may be generated at any time during the first 64 seconds after entry into holdover. Any such transient must be bounded by the MTIE mask in Figure 5-14, measured relative to a signal with the frequency of the input reference immediately before the reference loss.

- R5-123** [922] Any transient associated with entry into holdover shall be bounded by the MTIE mask in Figure 5-14.

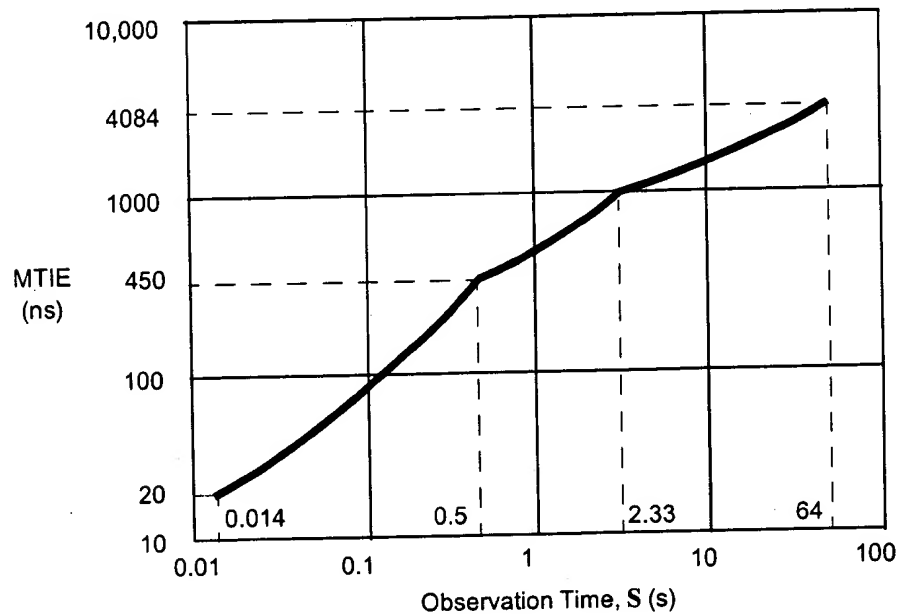
In addition to limiting phase transients during entry into holdover, requirements also exist to limit the fractional frequency offset (relative to the reference frequency immediately before the reference loss) that may occur during holdover.

- R5-124** [923] The initial fractional frequency offset, as defined in T1.105.09, shall be less than 0.05 ppm.
- R5-125** [924] The frequency drift rate, as defined in T1.105.09, shall be less than 5.8×10^{-6} ppm/second.
- R5-126** [925] The fractional frequency offset under varying temperature conditions shall not exceed 4.1 ppm.

Note that the first two of the requirements listed above are applicable independent of any temperature changes. However, they are tested against at constant temperature so that the components of the fractional frequency offset that they are intended to limit can be separated from any component caused by changing temperatures. The cumulative result of all three of these requirements is to limit the maximum holdover frequency offset to less than 4.6 ppm for the first 24-hour period of holdover under all operating conditions.

R5-127 [250] Entry into holdover, and restoration from holdover, shall be error-free.

Note that the preceding requirement does not apply to the traffic being carried on the failed reference. That traffic will either be lost (e.g., in a system with APS where both the working and protection lines fail) or may be temporarily interrupted while a protection switch is being completed (e.g., in a unidirectional path switched ring).



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.014$	N/A
$0.014 < S < 0.5$	$7.6 + 885 \times S$
$0.5 < S < 2.33$	$300 + 300 \times S$
$2.33 < S < 64$	$884 + 50 \times S$
$64 < S$	N/A

Figure 5-14. Phase-Transient for Entry into Holdover

5.4.4.2.3 Pull-in/Hold-in for SMCs

Pull-in and hold-in are defined in GR-1244-CORE. Pull-in/hold-in criteria are critical in maintaining the stratum hierarchy because they assure that any clock of a given stratum level will be able to take timing from any other clock of the same or higher stratum level.

- R5-128** [253] If a SONET NE with an SMC is timed from an external reference, the NE clock shall pull-in and hold-in to an external reference that is off frequency by ± 4.6 ppm (i.e., from a free-running stratum 3 clock).
- R5-129** [254] If a SONET NE with an SMC is timed from an OC-N reference, the NE clock shall pull-in and hold-in to an OC-N that is off frequency by ± 20 ppm.

It is important to realize that payload integrity is not necessarily guaranteed past frequency offsets of 4.6 ppm.

In environments where some NEs with SMCs and some NEs with stratum 3 internal clocks are deployed, possible interoperability issues exist. Specifically, an NE with a stratum 3 clock will not be able to pull-in to a signal from an upstream SMC in free-run. Careful engineering, following the hierarchical rules in GR-436-CORE, is necessary in architectures that mix NEs with SMCs and stratified clocks.

Conformance with pull-in/hold-in requirements is verified by testing for wander generation as per Section 5.4.4.3.2 after some "settling time." Note that during this settling time the NE may not conform with all wander generation and transfer criteria. Also note that this requirement is intended to apply in cases where the NE and clock are "warmed up" (e.g., after the frequency of the active reference changes). It is not intended to apply to clock start-up situations such as recovery after a power outage, or after the clock circuit pack is (re)inserted into the NE. That issue (i.e., clock warm-up times) may be discussed in a future issue of GR-253-ILR.

- R5-130** [926] The maximum settling time for an SMC shall be 100 seconds.

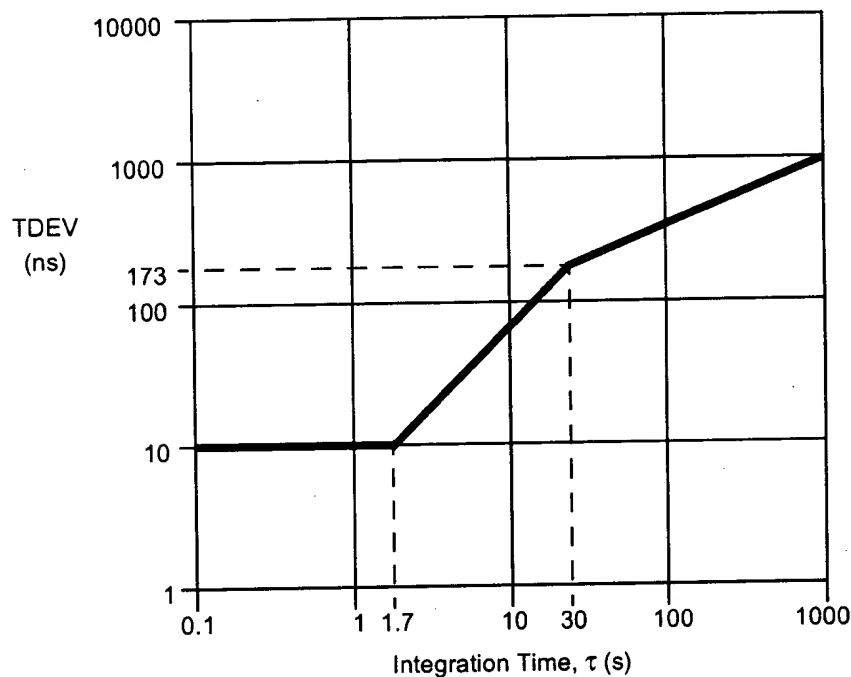
5.4.4.2.4 Wander Transfer for SMCs and Stratum 3 Clocks

The following criteria define the amount of phase noise filtering required of SMCs and stratum 3 clocks in SONET NEs. These requirements are based on work done in T1X1.3 (e.g., to control the amount of jitter on encoded DS3 payloads during worst-case external DS1 timing reference phase transients, and to limit the severity of pointer adjustment bursts).

- R5-131** [255v2] OC-N/OC-M and STS-N/STS-M electrical outputs, when referenced to an external timing signal that meets the wander TDEV mask

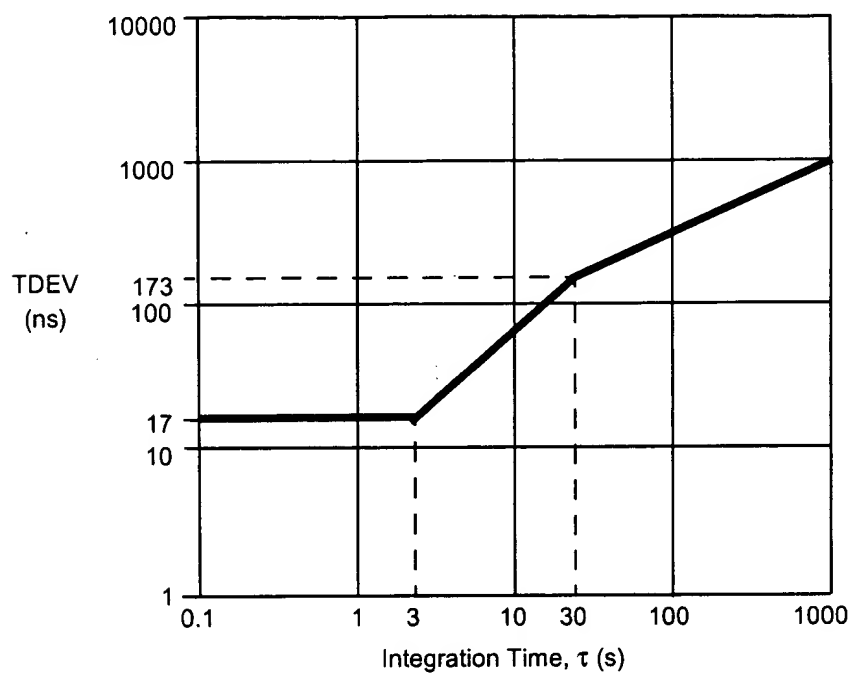
in Figure 5-16, shall be "less than or equal to" (where "less than or equal to" is defined to allow a phase gain of up to 0.2 dB in the pass-band of the clock) the wander TDEV mask given in Figure 5-15.

- R5-132** [256v2] OC-N/OC-M and STS-N/STS-M electrical outputs, when referenced to an OC-N timing signal that meets the wander TDEV mask in Figure 5-15, shall be "less than or equal to" (where "less than or equal to" is defined to allow a phase gain of up to 0.2 dB in the pass-band of the clock) the wander TDEV mask in Figure 5-15.



Integration Time, τ (seconds)	TDEV (nanoseconds)
$\tau < 0.1$	N/A
$0.1 < \tau < 1.7$	10
$1.7 < \tau < 30$	$5.88 \times \tau$
$30 < \tau < 1000$	$31.62 \times \tau^{0.5}$
$1000 < \tau$	N/A

Figure 5-15. OC-N Output Wander Time Deviation



Integration Time, τ (seconds)	TDEV (nanoseconds)
$\tau < 0.1$	N/A
$0.1 < \tau < 3$	17
$3 < \tau < 30$	$5.67 \times \tau$
$30 < \tau < 1000$	$31.62 \times \tau^{0.5}$
$1000 < \tau$	N/A

Figure 5-16. Time Deviation of Filtered Network Input to SONET NEs

5.4.4.3 All SONET Clocks

The following criteria apply to the clocks in all SONET NEs that contain LTE, independent of whether the clocks are stratum clocks or SMCs.

5.4.4.3.1 Clock Hardware

In general, large amounts of traffic are dependant upon the availability and quality of the clock. Therefore the clock in a SONET NE should be very reliable. To provide this level of reliability, the clock is required to be effectively duplicated by methods which allow corrective maintenance actions to be performed on failed elements without affecting the in-service elements.

- R5-133** [264v2] A stratum or SMC clock in a SONET NE shall meet the duplex equipment criteria in Section 3.3 of GR-1244-CORE.

5.4.4.3.2 Wander Generation

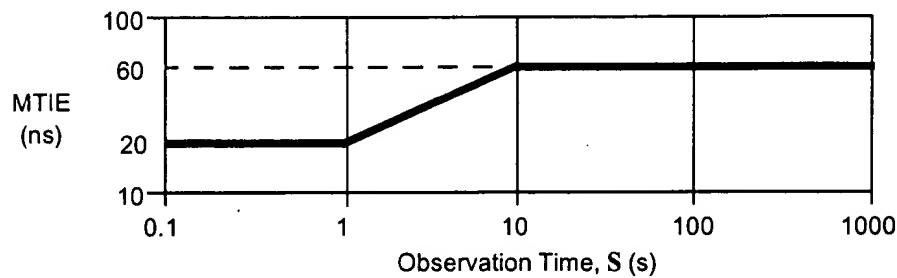
The intent of the following requirements is not to specify the amount of filtering that a SONET clock may provide, but rather to limit the amount of wander that the clock generates.

- R5-134** [257] OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the MTIE wander mask in Figure 5-17 when timed with a wander-free reference.
- R5-135** [258] OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the TDEV wander mask in Figure 5-18 when timed with a wander-free reference.

Conformance to these requirements is tested by providing an external or OC-N reference with bandlimited white noise phase modulation (i.e., jitter) of 1 μ s peak-to-peak. The jitter is bandlimited with 3-dB cutoffs at 10 Hz and 150 Hz.

The external timing input test signal described above was developed using ANSI Technical Report #6, *A Technical Report on Slave Stratum Clock Performance Measurement Guidelines*. That report discusses reference simulations with a 10 to 150 Hz filter and an amplitude of 1.5 μ s. Bellcore felt that an amplitude of 1.5 μ s was excessive, and therefore lowered it to 1 μ s in the Bellcore requirements documents. In addition, it has been suggested that this Technical Report is out of date, and that wander generation should be tested with a wander and jitter free signal. However, such a "clean" signal would not adequately stress the clock's timing recovery features,¹⁹ and therefore the timing test signal description given above is not expected to be changed.

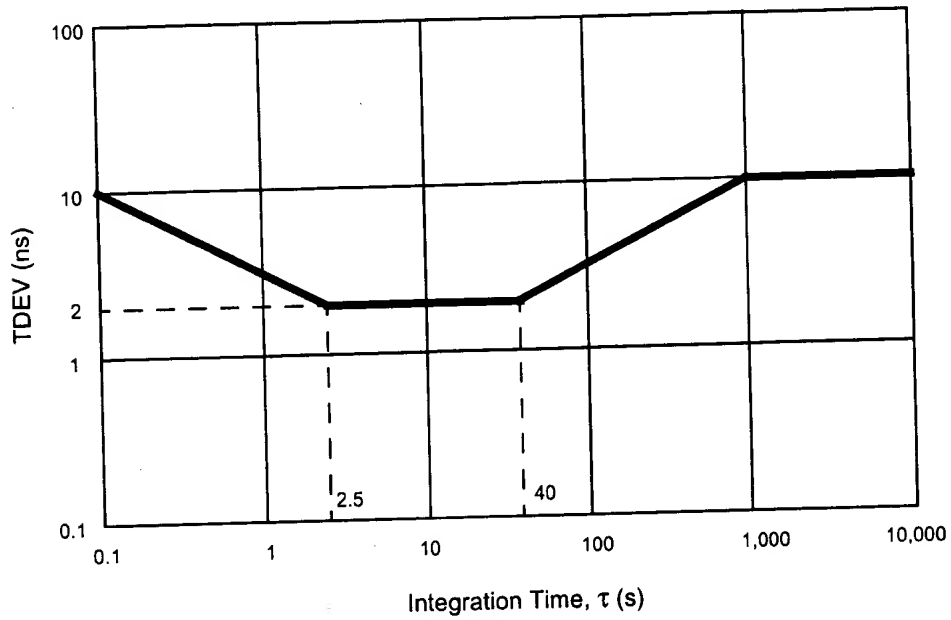
Note that the wander generation requirements are specified for arbitrarily long observation and integration times. The intent is that phase variations between the input and the output of the clock should be bounded. However, it is impractical to measure for infinitely long periods. It is recommended that compliance be verified using measurement periods up to 100,000 seconds.



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 < S < 1$	20
$1 < S < 10$	$20 \times S^{0.48}$
$10 < S$	60

Figure 5-17. MTIE for SONET Clocks

19. Note that a primary reason for using a test signal that has some noise (i.e., jitter) on it is to make sure that transitions occur in any digital components in an NE's timing recovery and signal generation chain. It is likely that such transitions will not occur during a test where the input signal is clean, but they will definitely occur (although possibly very infrequently) during normal network operations. Therefore, it is considered more important for the test signal to cause some transitions than it is for it to be a realistic example of a signal that could occur in the network.



Integration Time, τ (seconds)	TDEV (nanoseconds)
$0.1 < \tau < 2.5$	$3.2 \times \tau^{-0.5}$
$2.5 < \tau < 40$	2
$40 < \tau < 1000$	$0.32 \times \tau^{0.5}$
$1000 < \tau$	10

Figure 5-18. Time Deviation for SONET Clocks

5.4.4.3.3 Phase Transients

Phase transient criteria are important to control jitter on asynchronously mapped payloads, and also to control the generation of pointer adjustment bursts. Note that the criteria in this section that are applicable to SONET NEs containing stratum clocks are similar, but not identical, to the phase transient criteria in Issue 1 of GR-1244-CORE for stratum clocks used in other applications. In general, it is expected that GR-1244-CORE will be updated such that it will be consistent with the criteria that appear below (and with ANSI T1.101).

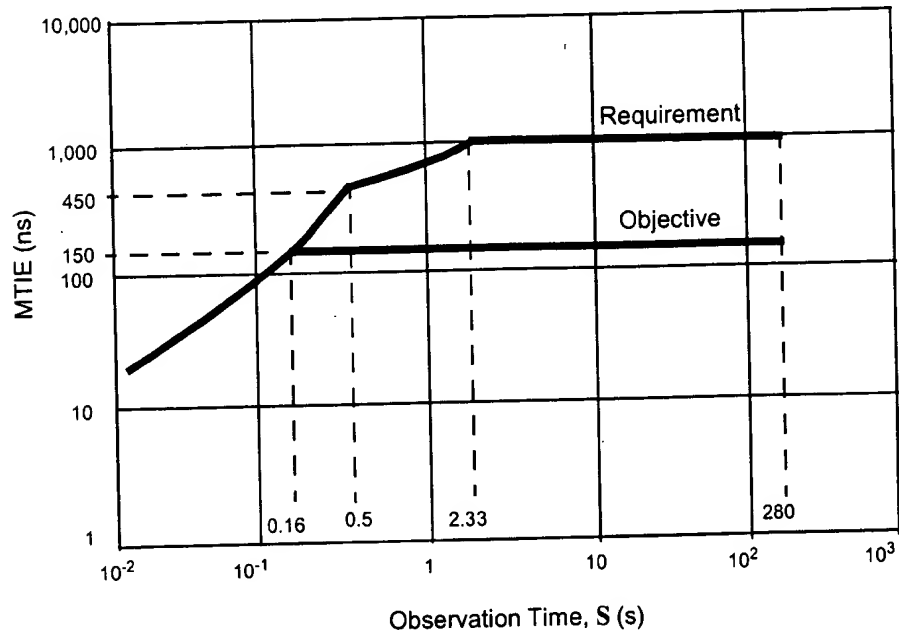
If that is done, this document may be modified to reference the phase transient criteria in GR-1244-CORE (wherever possible), rather than to provide duplicate criteria.

- R5-136** [259v2] For all SONET NEs that contain LTE, OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the specifications in ANSI T1.101-1994 for OC-N phase transients during synchronization rearrangement operations. Those specifications specify an MTIE of no greater than the "requirement" mask in Figure 5-19. Rearrangement activities include the following:
- ... • Manual timing reference switching
 - ... • Automatic timing reference switching as described in Section 5.4.6
 - ... • Switching between working line 1 and the protection line at the active reference OC-N "interface" (for NEs that support line APS)
 - ... • Entry into self-timing operation (i.e., holdover or free-run) for the initial 2.33 seconds of self-timing
 - ... • Automatic clock diagnostics
 - ... • Clock hardware protection switching.
 - ... • Phase transients on an external or OC-N synchronization input with the rate of change as specified in ANSI T1.101-1994

Note that although the same phase transient criteria apply, when an NE that is timed from an OC-N "interface" stops taking timing from the incoming signal on one line (e.g., working line 1) and begins to take timing from the signal on another line at the same "interface" (e.g., the protection line), it is not considered to have performed a timing reference switch (see Section 5.4.3.2).

It is desirable that phase transients from SONET NEs be small to minimize the number of STS pointer adjustments generated. The following criteria would allow for phase hits that cause no more than one STS pointer adjustment.

- R5-137** [1014] For SONET NEs that contain stratum 2 internal clocks, the MTIE of the SONET outputs during the internal rearrangements listed above (i.e., the first six rearrangements listed) shall meet the "objective" mask in Figure 5-19.
- O5-138** [260] For SONET NEs that contain stratum 3E, stratum 3 or SMC internal clocks, the MTIE of the SONET outputs during phase transients caused by the internal rearrangements listed above should be no greater than the "objective" mask in Figure 5-19.



Observation Time, S (seconds)	"Requirement" MTIE (nanoseconds)	"Objective" MTIE (nanoseconds)
$S < 0.0140$	N/A	N/A
$0.014 < S < 0.16$	$7.6 + 885 \times S$	$7.6 + 885 \times S$
$0.16 < S < 0.5$	$7.6 + 885 \times S$	150
$0.5 < S < 2.33$	$300 + 300 \times S$	150
$2.33 < S < 280$	1000	150

Figure 5-19. MTIE for Phase Transients from SONET Clocks

5.4.4.3.4 *Jitter and Errors During Synchronization Rearrangement Operations*

In general, the MTIE criteria in Section 5.4.4.3.3 could be interpreted to allow nearly instantaneous phase jumps of up to 20 ns. This is clearly undesirable from a jitter point of view. The following objective addresses this concern. It is expected that this objective will become a requirement in the future.

- O5-139** [928] The SONET outputs of an NE should meet the jitter generation requirement in Section 5.6.2.3.6 during the synchronization rearrangement activities listed in Section 5.4.4.3.3.
- R5-140** [261v2] Except for clock hardware protection switching, the synchronization rearrangement activities listed in Section 5.4.4.3.3 shall cause no errors on payload traffic.
- O5-141** [262v2] Clock hardware protection switching should cause no errors on payload traffic.

5.4.4.3.5 *Transition From Self-Timing to Normal Mode*

If an NE is in a self-timing mode (i.e., holdover or free-run) and a "good" reference becomes available, the user would normally want the NE to use that reference as the synchronization source. However, in some trouble-shooting and maintenance situations it may be appropriate for an NE with a stratum clock or SMC to continue self-timing. Therefore, the following criteria apply:

- R5-142** [265] Recovery from self-timing (holdover or free-run) shall be automatic.
- CR5-143** [266] An NE with a stratum clock or SMC may be required to provide the ability to inhibit automatic recovery from self-timing.
- R5-144** [267v2] Automatic restoration from holdover shall conform to the criteria in GR-1244-CORE, Sections 3.6 and 3.7. An SMC shall conform with the criteria for a stratum 3 clock.
- R5-145** [268] Automatic restoration from the free-run mode shall occur within two seconds of the presence of a validated reference signal.

In addition, a limit on the maximum rate of change of frequency during a stratum clock's or SMC's recovery from holdover is necessary to avoid the creation of excessive amounts of jitter on the SONET asynchronous payloads.

- R5-146** [930v2] The maximum rate of frequency change during holdover recovery shall be less than 2.9 ppm/second.

Measurement methodologies for testing an NE to these requirements are described in T1.105.09.

5.4.4.3.6 *Input Tolerance*

GR-1244-CORE (R1244-30) allows suppliers flexibility in determining when an NE will consider a reference failed. The NE may consider a reference failed as soon as it detects an LOS, AIS, OOF or LOF defect, or may it wait several seconds up to the point when a failure is declared, to see if the defect persists. If the NE does wait to see if the defect persists, its output synchronization performance must not be degraded.

- R5-147** [271v2] For interruptions (i.e., short LOS, AIS, OOF or LOF defects, not phase or frequency transients) of reference signals that do not cause reference switches or switches between lines (at an NE that supports line APS), the output criteria of Figure 5-17 shall be met.
- R5-148** [272] The NE shall tolerate phase transients on external and OC-N reference signals of a magnitude and slope as defined in ANSI T1.101-1994.
- R5-149** [1015] Clocks synchronized to an external DS1 timing signal shall tolerate, as a minimum, the jitter specified for the input test signal in the wander generation requirements in Section 5.4.4.3.2.
- CR5-150** [1016] Clocks synchronized to an external DS1 timing signal may be required to tolerate, as a minimum, input jitter applied according to the mask in Figure 7-1 of GR-499-CORE.
- R5-151** [1017] Clocks that are lined-timed from an incoming OC-N signal shall meet the Category II jitter tolerance requirements in Figure 5-28.

The term "tolerate" used in the previous criteria is defined to mean no errors on payload signals, no indication of improper operation (i.e., alarms), no reference rejection, and remaining frequency-locked to the reference so that the output phase variations, relative to the input reference, are bounded.

5.4.5 *Timing Distribution*

This section describes SONET NE criteria for SONET-based synchronization distribution. SONET-based synchronization distribution can occur two ways, either with DS1 signals

derived from a terminating OC-N,²⁰ or with retimed traffic DS1 signals carried in the SONET payload. The derived DS1 is the preferred method for timing distribution. Retimed payload DS1s are subject to the possibility of slips at the retiming buffer and should only be used for special applications. For example, for remote locations it may be preferable to provide a retimed signal rather than a separate timing-only signal that requires extra facilities.

5.4.5.1 Timing Distribution on Derived DS1 Signals

As stated in GR-436-CORE, "It is planned that synchronization distribution will evolve to become OC-N based." In order to realize this goal, the BITS clocks must be able to receive synchronization from SONET NEs. Since BITS clocks accept DS1 signals as input references, SONET NEs that contain LTE must be able to generate timing signals in the DS1 format.

- R5-152** [273v3] A SONET NE that contains LTE shall have the capability to supply two DS1 timing reference signals. The NE shall be capable of deriving both of these DS1s from a single line-side OC-N "interface" (see Figure 5-20) and, if more than one OC-N "interface" is supported, of deriving each DS1 from a different OC-N "interface" (see Figure 5-21). As a minimum, the derived DS1 signals shall be in the Superframe format and shall meet the ones-density requirement in GR-499-CORE.

The reason for specifying that both methods of deriving DS1 signals must be supported is that different applications have different uses for the derived DS1 signals. For example, if a SONET ring is being used for interoffice synchronization distribution and contains externally timed NEs that support DS1 synchronization status messages, then it would generally be beneficial to have each DS1 signal derived from a different OC-N "interface". Conversely, if the NEs in a SONET ring are line-timed or do not support DS1 synchronization status messages, then it would generally be beneficial to have both DS1 signals derived from the same OC-N "interface".

20. Although some NEs may support the capability to use an OC-M "interface" as a source for a derived DS1, that capability is not required in this document.

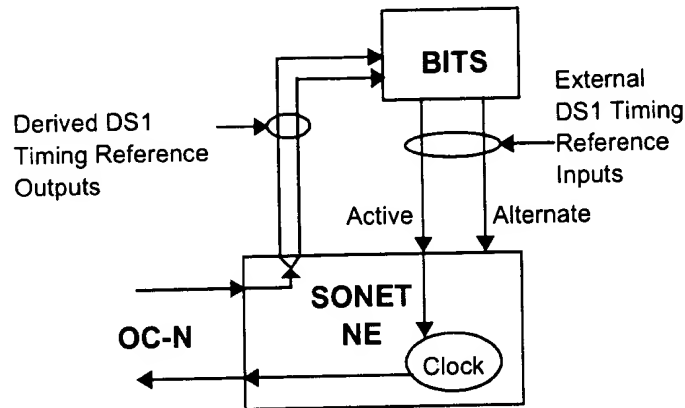


Figure 5-20. DS1 Timing References Derived from a Single OC-N "Interface" Example

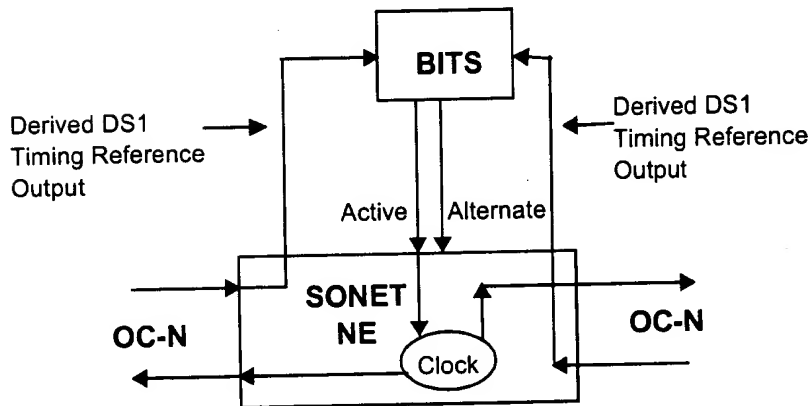


Figure 5-21. DS1 Timing References Derived from Different OC-N "Interfaces" Example

Note that in the case illustrated in Figure 5-20, if the NE supports line APS it may either (as a default) derive both DS1s from the OC-N signal on a single line (i.e., working line 1 or the protection line), or derive one DS1 from an OC-N on working line 1, and the other DS1 from the OC-N on the protection line. In any case, just as the reliability of SONET is enhanced by the ability to smoothly switch traffic to a protection line if a problem occurs

on a working line, the derived DS1 can use this protection switching capability to improve the robustness of the synchronization distribution network.

- CR5-153 [274v2]** An NE that supports line APS may be required to switch the source of timing for a derived DS1 to the OC-N on the protection line when the OC-N on working line 1 becomes unavailable as a derived DS1 source due to an LOS, LOF, or AIS (and vice versa if one of the DS1s is normally derived from the OC-N on the protection line).

In addition to switching the source of a derived DS1 between the working and protection lines at a single OC-N "interface", in some situations it may be desirable to support switching between separate OC-N "interfaces". Note that if this capability is supported, it implies that the capability is provided for the user to provision two (or possibly more) "interfaces" as potential sources for each derived DS1. Also note that if switching between "interfaces" is not desired in a particular application, it would normally be possible for the user to disable it (e.g., by provisioning only a single "interface" as the source for each derived DS1).

- CR5-154 [1018]** An NE that provides more than one OC-N "interface" may be required to support switching of the source of timing for a derived DS1 to a different (secondary) OC-N "interface" when the OC-N signal(s) at the original (primary) "interface" becomes unavailable as a derived DS1 source due to an LOS, LOF, or AIS, or when switching is appropriate based on the received synchronization status messages (see **R5-169 [285v3]**).

- R5-155 [1019]** If an NE that provides line APS supports switching between OC-N "interfaces" as the source of timing for a derived DS1, then it shall conform to **CR5-153 [274v2]**. In addition, a switch between "interfaces" in response to an LOS, LOF, or AIS shall occur only if the OC-N signals on both working line 1 and the protection line have failed.

As discussed in Section 5.4.5.2.1, switching the source of a derived DS1 from one OC-N "interface" to another at an externally timed NE is not recommended. Thus, any NE that is provisioned to support derived DS1 source switching is expected to also be provisioned to be line (or possibly through) timed. In addition, it is generally assumed that the same "interfaces" will be provisioned as potential timing references and potential derived DS1 sources, and that it is desirable for an "interface" that is currently being used as a timing reference to also be active as a derived DS1 source. To maintain this parallel use of OC-N "interfaces", the following requirements apply:

- R5-156 [1020]** A line-timed NE that is provisioned to support switching between OC-N "interfaces" as the source of timing for a derived DS1, and to (as a default) derive all of its active DS1s from the same OC-N "interface",

shall use the same type of switching (i.e., nonrevertive or revertive) as it uses for timing reference switching.

- R5-157** [1021] A through-timed ADM that is provisioned to support switching between OC-N "interfaces" as the source of timing for a derived DS1, and to (as a default) derive each DS1 from a different OC-N "interface", shall use revertive switching.

Note that if an NE's timing reference and derived DS1 source lists are identical, then maintaining the same active/inactive status for those two uses of a particular OC-N "interface" could be achieved by simply linking the derived DS1 source switching feature to the NE's timing reference and mode switching (i.e., through-timing to line-timing) features. Possible methods of automatically ensuring identical timing reference and derived DS1 source lists have been examined (see GR-253-ILR, Issue ID 253-51); however, they were not considered feasible. In addition, in some applications the user may want the lists to be different (at least temporarily). Therefore, it is left to the user to provision each list appropriately.

Also note that since switching between OC-N "interfaces" as the source for a derived DS1 is not recommended at externally timed NEs (see Section 5.4.5.2.1), no criteria have been included in this document to indicate if such switching (if supported) should be revertive or nonrevertive. Similarly, no revertive/nonrevertive switching criteria have been included to cover situations where parallel use of OC-N interfaces for timing and deriving DS1s is not desired (i.e., where a line-timed NE has been provisioned so that it normally derives its DS1s from the incoming signals at two different OC-N "interfaces", or where a through-timed ADM has been provisioned to normally derive both of its DS1s from a single OC-N interface).

In order to maintain the parallel use of an NE's OC-N "interfaces" for timing reference and derived DS1 source purposes in cases where the timing reference-related commands defined in Section 5.4.6 are used, it would be necessary for the NE to support the equivalent commands (e.g., a Manual Derived DS1 Source Switch command) for switching or locking out derived DS1 sources.

- CR5-158** [1052] A SONET NE that supports the capability to switch between OC-N "interfaces" as the source for a derived DS1 may be required to support derived DS1 source switching-related commands equivalent to the timing reference switching-related commands (defined in Section 5.4.6) that it supports.

- R5-159** [1053] If a SONET NE supports one or more derived DS1 source switching-related commands, the effects of those commands shall be functionally equivalent to the effects of the corresponding timing reference switching-related commands.

For example, a Manual Derived DS1 Source Switch command would need to be denied if it would cause a switch to a failed source or a source with a lower synchronization status message, while the effect of a Lockout a Derived DS1 Source command would be to effectively remove one source from the provisioned source list until the corresponding clear command is entered. (Note that similar to the case for the Forced Reference Switch command, the functionality of a Forced Derived DS1 Source Switch command is currently under study, see GR-253-ILR Issue ID 253-137.)

- O5-160** [276] The derived DS1 should be a framed all-ones signal.
- CR5-161** [277] The NE may be required to provide the capability for the user to provision the derived DS1 in the ESF format (in addition to the required Superframe format).
- R5-162** [278] The SONET NE shall be capable of supporting all its timing modes when providing derived DS1s from an OC-N.
- R5-163** [279v2] DS1 AIS (unframed all-ones; not the A-bit code for ESF) shall be inserted into the derived DS1 when the OC-N signal becomes unavailable as a derived DS1 source due to an LOS, LOF, or AIS (assuming no switching, see **CR5-153** [274v2], **CR5-154** [1018], **R5-155** [1019] and **R5-169** [285v3]). DS1 AIS shall be generated no later than the declaration of failure (see Section 6.2.1).
- R5-164** [280v2] Automatic restoration of the derived DS1 shall occur within two seconds of the clearing of the derived DS1 source failure.

In order to interface with BITS clocks and CPE, the derived DS1 signals have several performance criteria. It is expected that the derived DS1 from the SONET NE will meet ANSI T1.101 specifications. Some of the following performance criteria are more stringent than those in T1.101. Conforming with the following requirements will ensure that the derived DS1 signal, which may be generated at the end of a chain of clocks, meets the T1.101 specifications.

- R5-165** [281] The MTIE of the derived DS1 shall be less than 50 ns for observation times beyond the jitter region (observation times longer than 0.1 second).
- R5-166** [282] The TDEV of the derived DS1 shall be below the mask in Figure 5-22.

Currently, the conformance of an NE to these requirements is measured using a 10-Hz first-order low-pass filter, relative to an ideal, jitter-free and wander-free source that is also providing timing to the OC-N. Note however, that the characteristics of this source signal are under review, as discussed in GR-253-ILR, Issue ID 253-53. Also note that the MTIE

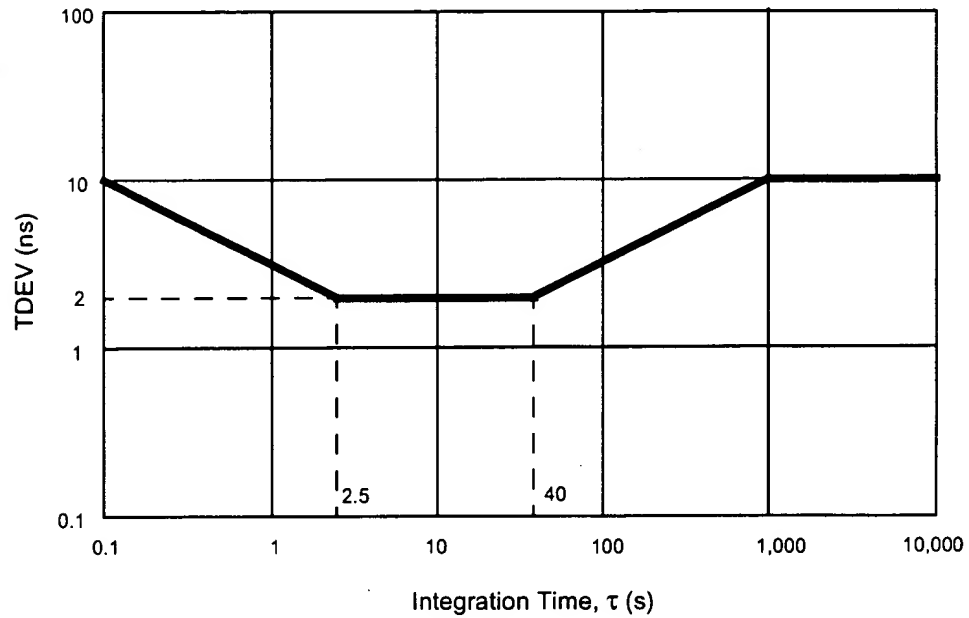
requirement does not specifically prohibit filtering incoming jitter from the OC-N, while the intent of the TDEV requirement is to force any wander that is generated on the derived DS1 to be at higher frequencies (lower integration times in terms of TDEV) so that it can be filtered by downstream BITS clocks.

R5-167 [283] The jitter on the derived DS1 shall be less than $1.0 U_{I_{pp}}$.

In this case, conformance testing is performed using a jitter-free OC-N source signal that meets the OC-N output wander TDEV mask in Figure 5-15 in Section 5.4.4.2.4. [For the purposes of this testing, an OC-N source signal can be considered to be "jitter-free" if the energy at observation times shorter than 0.1 s is such that the measured TDEV (in ns) is less than $100 \times \tau$ (with τ in seconds).]

R5-168 [284v2] The derived DS1 shall meet the MTIE requirement for DS1 phase transients in ANSI T1.101 during rearrangements (e.g., switching the derived DS1 source between working line 1 and the protection line in a system that supports line APS, or between "east" and "west" OC-Ns as per Section 5.4.5.2.1). For observation periods up to 280 seconds, a phase transient on a derived DS1 shall not exceed a magnitude of 1 μ s with a slope of 81 ns for a measurement period of 1.326 ms.

This measurement is made with a 100-Hz first-order low-pass filter, relative to an ideal, jitter-free and wander-free source.



Integration Time, τ (seconds)	TDEV (nanoseconds)
$0.1 < \tau < 2.5$	$3.2 \times \tau^{-0.5}$
$2.5 < \tau < 40$	2
$40 < \tau < 1000$	$0.32 \times \tau^{0.5}$
$1000 < \tau$	10

Figure 5-22. Time Deviation for Derived DS1

5.4.5.2 Synchronization Status Messages for Derived DS1 Signals

This section describes and contains the criteria related to the use of synchronization status messages on derived DS1 signals.

5.4.5.2.1 *Switching*

Ideally, the derived DS1 signals should always be referenced to the OC-N with the highest traceability as indicated by the synchronization status messages. However, switching between OC-N "interfaces" used to derive a DS1 at an externally timed NE (i.e., for use in the interoffice timing distribution network) would cause the maintenance and administration of the network, especially the avoidance of timing loop creation, to be extremely difficult. Conversely, switching between OC-N "interfaces" used to derive a DS1 at line-timed NEs and through-timed ADMs will not cause timing loops. Therefore, automatic switching of the OC-N "interfaces" used to derive the DS1 signals (in order to achieve the highest traceability) is viable only at line-timed NEs and through-timed ADMs.

- R5-169** [285v3] A line-timed NE or through-timed ADM shall automatically select (from the OC-N "interfaces" provisioned as possible sources for each derived DS1) the OC-N "interface" with the highest quality synchronization status message as the source for that derived DS1 signal.

Note that the preceding requirement covers synchronization status message-based source switching between OC-N "interfaces", but does not include switching between the signals on the working and protection lines at a single OC-N "interface". In general, working/protection line switching based on synchronization status messages is not required to be supported (although it may be provided by some NEs). The reason for this is that any NE that supports line APS is required to transmit the same synchronization status messages on working line 1 and the protection line (at a single "interface") in all situations. Thus, the messages that are received by the far-end NE on those two lines can be expected to be identical, and it is considered sufficient for that NE to monitor the incoming messages on only one of those lines (e.g., the line whose signal is currently being used as the derived DS1 source).

5.4.5.2.2 *Message Translation*

Two modes of translating synchronization status messages from the terminating OC-N to the derived DS1 outputs are defined for SONET NEs:

1. "threshold AIS generation" mode to be used when the receiving BITS clocks do not have synchronization status messaging capabilities

2. "message pass-through" mode to be used when BITS clocks have synchronization status messaging capabilities.

"Threshold AIS generation" mode is similar to the functionality currently provided by NEs that do not support synchronization messaging. Such NEs generate an AIS on the derived DS1 when the OC-N becomes unacceptable as a reference (e.g., during an LOS or AIS on the OC-N). Since BITS clocks recognize AIS as an unacceptable reference, the generation of AIS ensures that the BITS does not take synchronization from an OC-N that has been disqualified. With the "threshold AIS generation" mode, OC-N signals that degrade in terms of synchronization traceability also cause AIS, ensuring that the reference is rejected by the BITS. This mode allows the derived DS1 signal to be used as a synchronization reference in the cases where synchronization messaging in the DS1 ESF format is not supported by the BITS clock.

"Message pass-through" mode is appropriate in cases where the BITS clock supports synchronization status messaging. In this mode, the message on the OC-N is translated on to the derived DS1 and passed to the BITS. The BITS is then the element that decides which reference is appropriate, based upon the synchronization message. In order to support this mode, the DS1 signal must be in the ESF format.

Some service providers plan to upgrade their BITS clocks to support synchronization status messaging and therefore require the SONET NE to support the "message pass-through" mode. Other service providers do not plan to upgrade their BITS clocks and consequently do not require this capability.

R5-170 [286] An NE shall support the "threshold AIS generation" mode.

CR5-171 [287v2] An NE may be required by some service providers to support the "message pass-through" mode (i.e., it may be required to support synchronization status messages on the derived DS1).

R5-172 [288] If an NE supports both message translation modes, the format of the derived DS1 signal shall indicate which mode is used. If the derived DS1 is in the ESF format, the "message pass-through" mode shall be used. If the derived DS1 is in the SF format, the "threshold AIS generation" mode shall be used.

The purpose of the previous requirement is to minimize the amount of provisioning that is necessary. However, based on the requirement, if an ESF-formatted, derived DS1 signal is input to a BITS clock that does not support synchronization status messages, the BITS may select an unsuitable reference (e.g., a reference with a "DUS" message) to be active. Therefore, it is strongly recommended that the network provider use a derived DS1 in the ESF format only when the BITS supports synchronization status messages.

In "threshold AIS generation" mode, the following requirements apply:

- R5-173** [289] In "threshold AIS generation" mode, the NE shall insert AIS into the derived DS1 output when the synchronization status message in the OC-N signal that is being used as a reference for that derived DS1 is at or below a user selectable quality level. The default for the threshold shall be quality level 7, "SMC Traceable."
- R5-174** [290v2] The NE shall validate (as per Section 5.4.7.1) a degradation in the synchronization status message in the OC-N signal to or below the threshold, react as per Section 5.4.5.2.1, and generate AIS (if appropriate) on the derived DS1 timing output within 10 seconds.

In "message pass-through" mode, the following requirements apply:

- R5-175** [291] In "message pass-through" mode, the NE shall have the capability to generate the synchronization status messages listed in Table 5-7 in the ESF data link of the derived DS1 signals.
- R5-176** [292] The derived DS1 output shall carry the synchronization status message that corresponds to the message in the OC-N that is being used as a reference for that derived DS1.
- R5-177** [293v2] When the synchronization status message in the OC-N used to derive the DS1 changes, the NE shall validate the change (as per Section 5.4.7.1), react as per Section 5.4.5.2.1, and insert the appropriate message in the derived DS1, both within 10 seconds.
- R5-178** [294] The synchronization status message shall be sent continuously in the ESF data link.

5.4.5.3 Timing Distribution on Traffic-Carrying DS1 Payload Signals

This section provides criteria for a traffic-carrying DS1 signal that may be used for synchronization distribution to locations remote to the SONET NE. For timing distribution at locations where the SONET NE resides, the derived DS1 described in Section 5.4.5.1 is the preferred method. However, at remote locations it may be preferable to provide a signal retimed by a slip buffer, rather than a separate timing-only signal that would require extra facilities. DS1s that are byte-synchronously mapped may already be retimed by a slip buffer (see Section 3.4.1.1). If they are retimed, then they can be used for timing distribution purposes. However, if they are not retimed, or the NE only supports the asynchronous mapping, the following criteria apply:

- CR5-179** [295] An NE may be required to provide a retiming slip buffer for timing distribution on a traffic-carrying, payload DS1 interface.

The slip buffer retimes the output DS1 from the SONET NE's internal clock (which would be synchronized to an external reference or one of the NE's terminating OC-Ns). Bits are written into the buffer from the VT SPE, and are read out under the control of the SONET NE's clock. This buffer effectively smooths any phase transients due to VT pointer adjustments. It should be realized that the timing of the DS1 at the output of the slip buffer will be the same as the timing for the SONET NE. Therefore, it is important that the SONET NE have reliable timing. For example, in a line-timed access ring without synchronization status messages an NE with an SMC may end up in holdover during fiber cuts, which may cause unacceptable performance due to slips at the retiming buffer.

- R5-180** [296] The slip buffer shall be at least 1 frame (125 μ s) plus a minimum of 18 μ s of hysteresis. (More hysteresis is desirable.) When a slip occurs, an entire DS1 frame (i.e., 193 bits, including the framing bit) shall be slipped unless the DS1 is byte-synchronously mapped and the VT PTE is generating a new DS1 framing pattern. If a new DS1 framing pattern is being generated, then only the 192 data bits shall be slipped and the framing pattern shall not be affected.

Note that whenever the buffer slips for an asynchronously mapped DS1, the framing pattern is interrupted. This interruption may cause downstream reframes, but is considered necessary to provide a clear-channel service. These slip buffers are not considered a universal solution to the timing distribution problem, but may be useful in a few specific applications. It is critical that the signal being retimed is traceable to a PRS so that slips do not occur (except during failures).

- R5-181** [297] If a retiming slip buffer is provided, the NE shall accumulate slip counts as performance monitoring data according to the criteria in GR-820-CORE, *OTGR Section 5.1: Generic Digital Transmission Surveillance*.

5.4.6 SONET Timing Reference Switching

Section 3.4 of GR-1244-CORE contains various criteria related to timing reference switching. Except as noted below, those criteria are applicable to SONET NEs.

- R5-182** [298v3] A SONET NE shall meet the criteria related to automatic timing reference switching in Section 3.4 of GR-1244-CORE.

Note that several of the criteria referenced above, along with the criteria shown below related to manually initiated timing reference switches, refer to switching between references. Those particular criteria are applicable only to SONET NEs that support timing modes that utilize more than one possible reference (e.g., externally timed NEs, line-timed NEs with two OC-N "interfaces" provisioned as possible references).

- R5-183** [300v4] A SONET NE shall support a Manual Reference Switch command that allows the user to manually switch between synchronization references or to the specified reference. This command shall be denied if it would cause a switch to a failed reference or a reference with a lower synchronization status message (including a "DON'T USE for Synchronization" message). It shall also be denied if a Forced Reference Switch command is already active, or if the reference being switched to has been locked out via the Lockout a Reference command (if one or both of those commands are supported). In addition, a Manual Reference Switch shall be preempted if a change occurs or command is received that causes a subsequent timing reference switch (e.g., if the new active reference subsequently fails).

For an NE with only two provisioned references, it can be assumed that unless a particular reference is specified in the command, the intent of a Manual Reference Switch is to cause a switch from the currently active reference to the standby reference. On the other hand, if an NE has more than two provisioned references (or if the command specifies a particular reference at an NE that has only two provisioned references), then the important point is that after the command is completed, the specified reference is active.

Note that unlike the linear APS Manual Switch command defined in Section 5.3.6.1, an existing Manual Reference Switch command can be preempted by a subsequent Manual Reference Switch command (as well as by a reference failure or change in the incoming synchronization status messages on one or more of the references). Therefore it is not essential for another command to be defined to clear this command. However, such a command could be useful in some situations (e.g., if revertive reference switching is being used, it may be more intuitive for a user to enter a "Reference Switch Clear" command instead of a second Manual Reference Switch command to allow or cause the NE to switch back to the preferred reference). In addition, see the discussions and criteria below regarding additional commands related to reference switches that may be supported.

Also note that the Manual Reference Switch command defined above explicitly does not allow the NE to switch to a reference with a lower synchronization status message than the active reference, or to a failed reference. Thus, if (e.g., for maintenance purposes) the user wants the NE to use an available reference with a synchronization status message that is lower than the active reference, and if the command described above is the only synchronization reference switching command supported by the NE, then the user has to remove the active reference from the group of provisioned references. This allows the NE to switch to another reference during the maintenance period. When the work is complete, the user has to update the group of provisioned references to include the reference that was previously removed. Then the NE can automatically select the best reference to be active. To avoid the need for such reprovisioning (and the potential for errors that it introduces), the following additional reference switching-related commands are defined and may be implemented.

- CR5-184 [1054]** A SONET NE may be required to support a Forced Reference Switch command.

As discussed in GR-253-ILR Issue ID 253-137, certain details related to the functionality of the Forced Reference Switch command are currently under study. Therefore, if that command is supported by a particular SONET NE, it is important that its functionality be clearly documented (at least until those details are resolved and included in this document).

- R5-185 [1055]** If the Forced Reference Switch command is supported, the functionality of that command shall be clearly documented.

Although certain details of the Forced Reference Switch command functionality are still under study, various other details have been agreed upon. For example (and similar to the Manual Reference Switch case described above), if an NE's timing reference list contains only two possible references and the Forced Reference Switch command does not specify a particular reference, then the standby reference at the time the command is received is implied to be the specified reference. Also, an existing Forced Reference Switch command can be preempted by a subsequent Forced Reference Switch command, and a Clear command is definitely needed to allow the user to clear an existing Forced Reference Switch command.

- R5-186 [1056]** If the Forced Reference Switch command is supported, a Reference Switch Clear command to clear a Forced Reference Switch shall also be supported.

- CR5-187 [1057]** A SONET NE may be required to support a Lockout a Reference command.

- R5-188 [1058]** If the Lockout a Reference command is supported, its impact shall be to effectively cause the specified reference to be temporarily suspended from the NE's provisioned timing reference list (i.e., until the corresponding Clear command is received).

Note that if all of an NE's other provisioned references are considered failed or unavailable or contain the DUS synchronization status message when (or after) the Lockout a Reference command is performed, the NE would need to enter holdover according to the criteria in GR-1244-CORE.

Unlike the Manual and Forced Reference Switch command cases, a newly entered Lockout a Reference command does not preempt an existing Lockout a Reference command. In addition, its effect on a new or existing Manual or Forced Reference Switch command depends on the particular references specified in those commands. That is, a Manual or Forced Reference Switch command can coexist with Lockout a Reference command if it applies to a different reference, but is denied or preempted if it applies to the same reference as specified in the Lockout a Reference command. Thus, the Lockout a Reference command is analogous to the linear APS Lockout a Working Channel command defined in

Section 5.3.6.2, and can be used to effectively remove one or more specific references from the list of references at an NE with multiple provisioned references (until the corresponding clear command is entered).

- R5-189** [1059] If the Lockout a Reference command is supported, a Reference Lockout Clear command shall also be supported. That command shall cause the NE to clear an existing Lockout a Reference for the specified reference.

The following requirement is intended to cause an NE that supports the Forced Reference Switch command and/or the Lockout a Reference command to alert the user when those commands are active. Note that it is written such that only a single alarm or event message (and a single clear message) would be sent to the OS if, for example, the user entered a series of Forced Reference Switch commands with no Clear commands in between (i.e., with each Forced Reference Switch command preempting the previous one).

- R5-190** [1060] A SONET NE shall declare (and report to an OS) a standing condition when the first of one or more concurrent or consecutive Forced Reference Switch or Lockout a Reference commands is completed. The level of this condition shall be user-provisionable either as Not Alarmed or as a MN alarm. The condition shall be cleared (and a clear message sent to an OS) when all Forced Reference Switch and/or Lockout a Reference commands have been cleared.

5.4.6.1 Timing Reference Failure Conditions

In general, the conditions for a SONET NE to consider a timing reference failed are the same as those in Section 3.4.1 of GR-1244-CORE. Therefore the criteria in that section are applicable to SONET NEs (see **R5-182 [298v3]**), along with the additional criteria discussed below.

- R5-191** [1022] An incoming SONET signal shall be considered failed or unavailable for timing purposes under the following conditions:
- ... • Loss of signal energy (e.g., LOS defect detection/failure declaration)
 - ... • Loss of framing (e.g., LOF defect detection/failure declaration)
 - ... • Line AIS (e.g., AIS-L defect detection/failure declaration at an NE containing LTE)

Note that the preceding requirement is included here because the corresponding criteria in Issue 1 of GR-1244-CORE specifically address failures only on DS1 and CC timing references (i.e., they do not address failures on SONET signals). As discussed in GR-1244-ILR, Issue ID 1244-4 and GR-253-ILR, Issue ID 253-59, GR-1244-CORE is

expected to be updated to cover failures on OC-N signals that are used for timing, and at that time this requirement may be deleted. Also note that (consistent with the discussion in GR-1244-CORE) this requirement provides some flexibility with respect to the time that an NE takes to consider a reference failed. Specifically, an NE is not required to consider a signal to be failed for timing purposes immediately upon detecting a particular defect, and is also not required to delay considering the signal as failed until the corresponding failure is declared. What is important is that the NE conform to the applicable criteria independent of whether or not it considers the signal failed (e.g., **R5-147 [271v2]** in Section 5.4.4.3.6 in situations where it is not considered failed).

- R5-192 [302v2]** A timing reference shall be considered failed or unavailable upon receipt of synchronization status message indicating that the reference is traceable to a source in holdover/free-run that is of worse quality than the local internal clock.

For example, a SONET NE with an internal stratum 3 clock would consider a reference failed if it has a synchronization status message of either traceable to an SMC or a stratum 4 clock.

- R5-193 [305]** If the active timing reference is an OC-N "interface", switching to an alternate timing reference (if available) shall only take place after it has been determined that any available protection switching of the OC-N line and its terminating circuitry has failed to end the outage. The clock shall maintain accuracy and meet the phase transient criteria in Section 5.4.4.3.3 during the protection switch. Traffic and timing need not be protected together.

Note that in order for an NE that does not support provisioning of an OC-N/M "interface" as a single reference (see **R5-110 [1051]**) to meet **R5-193 [305]**, the working and protection lines would need to be listed consecutively in the NE's timing reference list. In addition, such an NE may use the same type of switching (i.e., revertive or nonrevertive) between references at a single "interface" as between references at separate "interfaces". (See GR-253-ILR, Issue ID 253-114 regarding the type of switching to be used for switches between lines at a single "interface" when that "interface" is considered a single reference.)

Also note that although the receipt of a "DON'T USE for Synchronization" message disqualifies a reference from being active, this condition does not cause the reference to be considered failed. This requirement may seem contrary to current alarm methods that define the unavailability of a redundant reference to be a minor alarm condition. However, alarming the receipt of a "DON'T USE for Synchronization" message could lead to standing alarms at SONET NEs. It is widely agreed that standing alarms should be avoided. Furthermore, the receipt of a single "DON'T USE for Synchronization" message is a normal condition in some applications. The purpose of generating an alarm is to alert people that something is abnormal and needs to be corrected.

As identified in Section 3.4.1 of GR-1244-CORE, a SONET NE may be required to reject a reference based on frequency offset by some service providers in some applications. In particular, some service providers are concerned that payload performance is not guaranteed for frequency offsets greater than 4.6 ppm, while clocks in SONET NEs may be pulled off-frequency to offsets greater than this. However, even stratum 3 clocks can not provide frequency rejection at 4.6 ppm.

5.4.6.2 Performance During Timing Reference Switching

Refer to Section 5.4.4.3.4 for the rearrangement phase transient limits that are applicable to reference switches. The MTIE requirement does not apply for reference switches caused by references that fail due to off-frequency conditions.

5.4.6.3 Revertive and Nonrevertive Timing Reference Switching

The applicable criteria for revertive and nonrevertive switching between references is in Section 3.4.3 of GR-1244-CORE (see **R5-182 [298v3]**). Note that those criteria are specifically intended for switching between references, not switching between the working and protection lines at a single OC-N "interface". The appropriate criteria for that case are under study (see GR-253-ILR, Issue ID 253-114).

5.4.6.4 Synchronization Status Messages and Timing Reference Switching

The following reference switching requirements apply for SONET NEs that support synchronization status messages.

R5-194 [310] The available, provisioned reference with the highest quality synchronization status message shall be selected as the active reference.

R5-195 [311] The NE shall not select a reference with a "DON'T USE for Synchronization" message as the active synchronization reference.

Note that the preceding requirements cover synchronization status message-based switching between references (e.g., OC-N "interfaces"), but do not include switching between the signals on the working and protection lines at a single OC-N "interface". In general, working/protection line switching based on synchronization status messages is not required to be supported (although it may be provided by some NEs). The reason for this is that based on the existing requirements, any NE that supports line APS should transmit the same synchronization status messages on working line 1 and the protection line (at a single "interface") in all situations. Thus, the messages that are received by the far-end NE on those two lines can be expected to be identical, and it is considered sufficient for that NE

to monitor the incoming messages on only one of those lines (e.g., the line whose signal is currently being used as the timing source).

Also note that requirement **R5-194 [310]** implies that when the active reference has a degradation in synchronization quality or an alternate reference has an improvement in synchronization quality, a reference switch may be triggered. In many cases, it would be appropriate to perform this switch "immediately". However in some cases, the absence of a "holdoff" time before the switch is performed will result in extraneous reference switches. In particular, this is expected to be an issue for externally timed NEs that are timed from a BITS that does not change the synchronization status messages on all of its outputs simultaneously.

O5-196 [1023] An externally timed NE should refrain for at least X seconds from performing a switch between its primary and secondary external DS1 reference signals based on a change in the synchronization status messages contained in those signals, unless that change causes the NE to consider its currently active reference as failed or unavailable (see **R5-192 [302v2]**), or the message in the active reference changes to the "DON'T USE for Synchronization" message (see **R5-195 [311]**). If a reference switch is still appropriate at the end of this X -second holdoff period, the NE should perform the switch within an additional 1-second period.

The current value of X is '10'. However, this value may need to be settable and is currently under study (see GR-253-ILR, Issue ID 253-60).²¹ Once the value is considered to be stable, it is expected that **O5-196 [1023]** will be changed to a requirement.

Table 5-8 contains an example in which the synchronization status messages on two references go through a series of changes and indicates which reference would be selected by the NE to be active. Two scenarios, revertive and nonrevertive reference switching, are illustrated. Note that the entries in the table are not static, individual examples. Instead, the table represents a series of events starting with the conditions of the first row, and going through various changes to the conditions of the last row.

21. The ability to set the value of X would be needed if future standards activities reduce the time that a BITS is allowed to finish changing the messages on its outputs.

Table 5-8. Example of Reference Selection Using Synchronization Messages

Sync Message Reference A	Sync Message Reference B	Active Reference Nonrevertive	Active Reference Revertive (default A)
PRS	STU	A	A
PRS	PRS	A	A
ST2	PRS	B	B
ST2	STU	B	B
ST2	PRS	B	B
PRS	PRS	B	A

5.4.7 Message Validation and Generation

This section discusses the synchronization messaging criteria for SONET NEs. The specific criteria describe how SONET NEs validate, react to, and generate synchronization status messages.

5.4.7.1 Message Validation

These criteria are based upon the functions specified in T1 Technical Report #33. It is important that the scheme used to validate the synchronization status messages be robust to occasional errors in the messaging channel.

- R5-197** [312] The supplier shall document the scheme used to validate messages on the ESF data link for DS1 synchronization references and on the S1 byte for SONET signals.
- R5-198** [313] A change in the S1 synchronization status message shall be detected if at least 8 consecutive samples (these may or may not be consecutive SONET frames) of bits 5-8 of the S1 byte have the same (new) value. The sampling rate shall be such that the maximum time to detect a change (assuming no transmission errors) is 1 second.
- R5-199** [314] A valid synchronization status message in the ESF data link for DS1 synchronization signals shall be detected if at least 7 out of 10 like messages are received.

The NE uses the format of the synchronization reference to determine how to react when the reference does not have a recognizable synchronization message. The NE behaves differently depending upon whether the reference is an OC-N or an external DS1 reference.

- R5-200** **[315v2]** For S1 messages, if no validated synchronization status message is detected (e.g., due to transmission errors or the receipt of an undefined message) for a period of greater than 10 seconds, the NE shall consider the reference failed.

If a DS1 reference is in the SF format, messaging cannot be expected.

- R5-201** **[316]** For DS1 references in the SF format, the SONET NE shall consider the reference to have a "Synchronized – Traceability Unknown" message.

If a DS1 reference is in the ESF format, the NE can initially assume that the reference supports synchronization status messaging. Therefore, the NE can expect to see a valid message on the reference. However, an external reference in the ESF format may sometimes be used without synchronization messages. In this case, the user would want the NE to behave differently. A provisioning step is necessary to make the distinction between an ESF signal that supports synchronization messages and one that does not.

- R5-202** **[318]** The user shall be able to provision the NE to accept an external DS1 reference in the ESF format that does not support synchronization messages. For such a provisioned reference, the SONET NE shall consider the reference to have a "Synchronized – Traceability Unknown" message.

- R5-203** **[317v2]** For DS1 references in the ESF format, if no validated synchronization status message is detected (e.g., due to transmission errors) for a period of greater than 10 seconds, the SONET NE shall consider the reference to be failed (unless the reference has been provisioned as not supporting synchronization status messages).

The user needs to be able to determine the synchronization message at different interfaces to the NE. In addition, the synchronization message at certain interfaces (e.g., those used as timing references) is important for network management so the NE should automatically report changes in the message to the OS.

- R5-204** **[319]** The NE shall automatically generate a status report to an OS when the synchronization status message on any provisioned reference changes. The report shall indicate which reference changed, the time of the change, the old synchronization status message, and the new synchronization status message.

- O5-205** [320v3] The NE should report to the user, on demand, the synchronization status message at any of its SONET "interfaces" (output and input, including those where processing of the incoming S1 byte has been disabled), and on its external DS1 reference signals (when supported).

5.4.7.2 Message Reaction

The criteria in this section apply to all NEs that support synchronization status messaging.

- R5-206** [321] When the synchronization status message in the active synchronization reference changes, the NE shall validate the change as per Section 5.4.7.1, react as per Sections 5.4.5.2.1 and 5.4.6.4, and insert the appropriate synchronization status message in all transmitted SONET signals²² within 10 seconds.
- R5-207** [322] When the NE enters holdover or free-run, the synchronization status message on all of its transmitted SONET signals shall be changed within 10 seconds to indicate the holdover level of the SONET NE's internal clock (e.g., "Stratum 3 Traceable" or "SMC Traceable").
- R5-208** [323] When the NE recovers from holdover or free-run, the synchronization status message shall not change until the NE has completely resynchronized. The time to change the message shall be no longer than the recovery time [i.e., the sum of the requalification time (as specified in Section 5.4.4.3.5) and the locking time (as specified in GR-1244-CORE)] plus 10 seconds.

5.4.7.3 Message Generation

A SONET NE generates the synchronization message in the outgoing SONET signals differently depending upon its timing mode. The criteria for externally timed, line-timed, and through-timed NEs are detailed in separate subsections. If the NE support synchronization messaging, the message in the S1 byte should be sent at all times except as noted in the following requirement.

- R5-209** [324] The NE shall allow the user to individually provision each SONET "interface" so that the transmitted signals from that "interface" carry the

22. "All transmitted signals" includes any SONET signal transmitted from an "interface" where synchronization status messages are supported. See Section 5.4.2 for the criteria on supporting the synchronization status messages at line-side and drop-side "interfaces".

“DON'T USE for Synchronization” message instead of the message that reflects the actual traceability of the signals.

The purpose of this requirement is to block the synchronization status messages on certain interfaces. Typical applications would be at carrier-to-carrier or carrier-to-customer interfaces. The service providers currently do not have control over whether the traffic signals are used as timing references. This requirement gives service providers the capability to block certain signals as synchronization references.

5.4.7.3.1 *Externally Timed NEs*

When available, external timing is the preferred timing mode for SONET NEs. These criteria are relevant for all externally timed NEs, particularly those in central offices.

In Section 5.4.7.1, it was described how the NE reacts differently depending upon the format of the external DS1 reference. If the external reference is in the SF format, the NE knows that messaging cannot be expected on the signal. Conversely, the NE assumes that an ESF DS1 reference supports synchronization messages. A similar distinction based on the format of the external DS1 reference is made for the message generation criteria.

If the external DS1 reference is in the ESF format, the following requirements apply (unless the NE has been provisioned to not expect messages on the external ESF DS1 reference).

- R5-210** [325] If none of the terminating signals at a particular SONET “interface” are being used to derive a DS1, then the NE shall insert the synchronization message from the active external ESF DS1 reference on the SONET signals transmitted from that “interface”.
- R5-211** [326v2] If a terminating signal at a particular SONET “interface” is being used to derive a DS1 and the synchronization status message on the active external ESF DS1 reference matches the synchronization status message on the terminating SONET signal, then the NE shall insert the “DON'T USE for Synchronization” message on the SONET signals transmitted from that “interface” (unless the automatic generation of the DUS message has been disabled for all of the DS1s being derived from that interface, see **CR5-214** [1061] and **R5-215** [1062]).
- R5-212** [327v3] If a terminating signal at a particular SONET “interface” is being used to derive a DS1 and (in the steady-state, see **R5-213** [1024]) the synchronization status message on the active external ESF DS1 reference does not match the synchronization status message on the terminating SONET signal (or the automatic generation of the DUS message has been disabled for all of the DS1s being derived from that interface, see **CR5-214** [1061] and **R5-215** [1062]), then the NE shall insert the synchronization

message from the active external reference on the SONET signals transmitted from that "interface".

The effect of these requirements is illustrated in Figure 5-23, in which OC-N #1 carries the DUS message in accordance with **R5-211 [326v2]**, and OC-N #2 carries the PRS message in accordance with **R5-212 [327v3]**.

- R5-213 [1024]** When an NE that is transmitting the DUS message at one of its SONET "interfaces" (because it meets the condition described in **R5-211 [326v2]**) detects a change in the incoming synchronization status message at that "interface", it shall continue to generate a DUS message for *Y* seconds after the synchronization status message has been translated to the derived DS1 (i.e., after the message on the outgoing derived DS1 has been changed to reflect the new message received at that SONET "interface").

The current value of *Y* is '10'. However, this value may need to be settable in case future standards activities reduce the time that a BITS is allowed to detect changes on its references and react appropriately (e.g., change the messages on its outputs). Once that time is considered to be stable, it is expected that *Y* will be replaced with a specific value. After the period of *Y* seconds, **R5-211 [326v2]** and **R5-212 [327v3]** are again the applicable criteria for determining the message to be generated on the transmitted SONET signals.

In certain applications (e.g., in applications where a particular derived DS1 is not being used as a reference signal for the BITS), it may be desirable for an externally timed NE to provide the user with the capability to disable the automatic generation of the DUS message according to the requirements shown above. To support such applications, the following criteria apply. Note that similar criteria related to disabling the automatic generation of the DUS message by a line-timed SONET NE do not exist, and that it is recommended that such a feature not be provided.

- CR5-214 [1061]** An externally timed SONET NE may be required to be capable of being provisioned such that it does not automatically generate the DUS message according to **R5-211 [326v2]**.
- R5-215 [1062]** A SONET NE that can be provisioned to derive each of its DS1s from a different SONET "interface" (see **R5-152 [273v3]**) and that also provides the capability to disable the automatic generation of the DUS message according to **R5-211 [326v2]**, shall provide that capability on a per-derived DS1 basis.
- R5-216 [1063]** If a SONET NE provides the capability to disable the automatic generation of the DUS message according to **R5-211 [326v2]**, its default shall be that the automatic generation is enabled.

If the external DS1 reference is in the SF format or the NE has been provisioned to not expect messages on the external ESF DS1 reference, the following requirement applies.

- R5-217** [328] If the external reference does not carry synchronization messages (e.g., an external DS1 reference is in the SF format or the NE has been provisioned not to expect synchronization messages on the ESF DS1), the NE shall insert the “Synchronized – Traceability Unknown” message on all transmitted SONET signals.

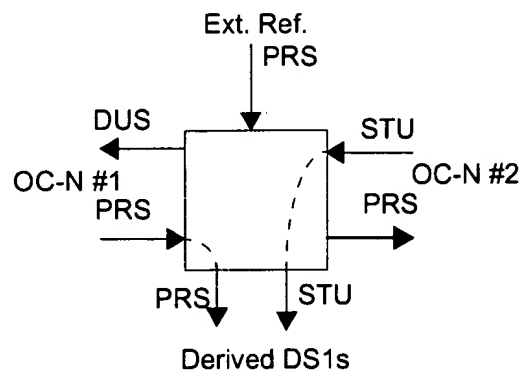


Figure 5-23. Example Implementation of R5-211 and R5-212

5.4.7.3.2 Line-timed NEs

When the NE is line-timed, the following requirements apply.

- R5-218** [329] At all SONET “interfaces” that are not the active synchronization reference, the line-timed NE shall insert the synchronization status message from the active synchronization source on the transmitted SONET signals (see Figure 5-24, OC-N #2).
- R5-219** [330] The line-timed NE shall generate a “DON’T USE for Synchronization” message in the transmitted signals from the OC-N “interface” that is the active synchronization reference (see Figure 5-24, OC-N #1).

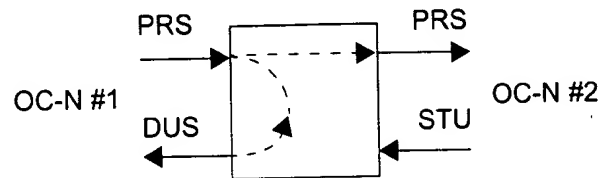


Figure 5-24. Example Implementation of R5-218 and R5-219

5.4.7.3.3 Through-Timed NEs

When a SONET ADM is through-timed, the following requirements apply.

R5-220 [331] A through-timed ADM shall insert the synchronization message from the terminating OC-N in the transmitted OC-N in the same direction (see Figure 5-25 for an example).

R5-221 [332] A through-timed ADM shall insert the synchronization status message of the appropriate timing source in all dropped SONET signals.

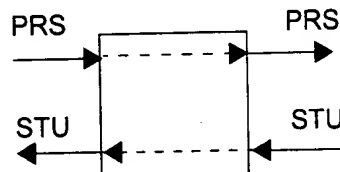


Figure 5-25. Example Implementation of R5-220

5.5 Framing For SONET Signals

This section contains the criteria related to the monitoring of incoming SONET signals for Severely Errored Framing (SEF) defects. These criteria are applicable to all SONET NEs with STE functionality, and also to non-STE NEs (e.g., physical layer regenerators) that need to frame on the SONET signal for monitoring purposes. Also see Section 6.2.1.1.2 for related criteria on Loss of Frame defects and failures (which occur if an SEF defect persists).

- R5-222** [333] The framing pattern observed by a SONET NE shall include a subset of the A1 and A2 bytes contained in the incoming STS-N electrical or OC-N signal.
- R5-223** [334] An SEF defect shall be detected when the incoming signal has a minimum of four consecutive errored framing patterns. The maximum SEF detection time shall be 625 μ s for a random signal.
- R5-224** [335] The framing algorithm used to check the alignment shall be such that an SEF defect is not detected more than an average of once every 6 minutes while the BER of the STS-N electrical or OC-N signal is 10^{-3} .²³
- R5-225** [336] Once an SEF defect has been detected, the SONET NE shall terminate the SEF defect upon detecting two successive error-free framing patterns.

Any implementation of the frame recovery circuitry that, following a detected SEF defect, achieves realignment within the 250- μ s interval implied by the preceding requirement is acceptable.

Note that it is not necessary for a SONET NE to use the same framing pattern for detecting SEF defects as it uses for terminating SEF defects (i.e., it may monitor different subsets of the A1 and A2 bytes for detecting and terminating the SEF defect). In addition, a framing pattern does not necessarily have to consist of full bytes. For example, an NE could monitor certain bits of a particular A1 byte, along with all or parts of one or more other A1 and A2 bytes for detecting SEF defects.

5.6 Jitter

This section provides jitter criteria for SONET transport systems and equipment. These criteria limit the timing jitter at SONET network connections.

Jitter is defined as the short-term variations of a digital signal's significant instants from their ideal positions in time.²⁴ Significant instants could be (for example) the optimum sampling instants. Long-term variations correspond to wander and are addressed as part of synchronization criteria in Section 5.4, and the phase variation criteria in Section 5.7.

23. Poisson distribution of bit errors is assumed.

24. "Short-term variations" implies phase oscillations of frequency greater than or equal to some demarcation frequency. Currently, 10 Hz is the demarcation between jitter and wander in the DS1 to DS3 North American Hierarchy. The demarcation frequencies for the SONET hierarchy correspond to the B_1 values in Table 5-9.

5.6.1 Network Interface Jitter Criteria

The following requirement limits the levels of jitter at OC-N and STS-N electrical interfaces. These limits restrict the broadband jitter appearing on a SONET signal anywhere in a SONET line system. Therefore, the requirement applies at interfaces between users and carriers, as well as between two carriers.

- R5-226** [337] Timing jitter at network interfaces shall not exceed A_1 Unit Intervals peak-to-peak (UI_{pp}) when measured over a 60-second interval with a bandpass filter having a high-pass cutoff frequency of B_1 (and a roll-off of 20 dB/decade) and a low-pass cutoff frequency of at least B_3 .²⁵
- ... Timing jitter at network interfaces shall not exceed $A_2 UI_{pp}$ when measured over a 60-second interval with a bandpass filter having a high-pass cutoff frequency B_2 (and a roll-off of 20 dB/decade), and a low-pass cutoff frequency of at least B_3 .

The values for the A_1 , A_2 , B_1 , B_2 , and B_3 parameters referred to in this requirement appear in Table 5-9. Parameter values are given only for OC-N and STS-N electrical interfaces where N is equal to 1, 3, 12, and 48. Parameter values for other values of N either do not exist, or appear in other documents (e.g., ANSI T1.105.03, *Synchronous Optical Network (SONET) - Jitter at Network Interfaces*, for OC-24, GR-1377-CORE for OC-192). Note that two sets of values are given for OC-48 interfaces. The second set [OC-48(B)] corresponds to a tighter jitter limit than the first set (i.e., the limit with $B_2 = 12$ kHz is more difficult to meet than the limit with $B_2 = 1000$ kHz). This tighter limit must be met to ensure proper performance in the presence of NEs exhibiting "reduced" jitter tolerance. See Section 5.6.2.2.2 for further information on reduced jitter tolerance.

25. The amount of jitter at frequencies above B_3 is expected to be very small. Therefore the low-pass cutoff frequency and roll-off characteristics of the bandpass filter used in jitter generation and network interface jitter measurements are not expected to have significant impacts on those measurements. Thus, no particular low-pass roll-off characteristics are specified, and the low-pass cutoff frequency is specified using the term "at least" (i.e., higher low-pass cutoff frequencies are also acceptable).

In addition, if there is a significant amount of jitter at high frequencies, that jitter could adversely impact the performance of the system and the user should know that it is present. Therefore, it is desirable to include as much of the high-frequency jitter as possible in the measurements.

Table 5-9. Parameters for Network Interface Jitter Requirements

OC-N/STS-N Level	B ₁ (Hz)	B ₂ (kHz)	B ₃ (MHz)	A ₁ (UI _{pp})	A ₂ (UI _{pp})
1	100	20	0.4	1.5	0.15
3	500	65	1.3	1.5	0.15
12	1000	250	5	1.5	0.15
48	5000	1000	20	1.5	0.15
48(B)	5000	12	20	1.5	0.15

To ensure that the accumulated jitter at an interface does not exceed the network limits, and that SONET line systems do not experience performance degradations due to jitter, the criteria in Section 5.6.2 are applicable to SONET NEs.

5.6.2 SONET NE Jitter Criteria

SONET equipment jitter criteria are specified in the following areas:

- Jitter transfer (Section 5.6.2.1)
- Jitter tolerance (Section 5.6.2.2)
- Jitter generation (Section 5.6.2.3).

The criteria in these sections specify limits for SONET equipment interfaces that fall into the following two categories (see GR-499-CORE for the definitions of these categories):

- Category I – Asynchronous DS_n interfaces to SONET NEs are considered Category I interfaces.
- Category II – OC-N, STS-N electrical, and synchronous²⁶ DS1 interfaces to SONET NEs are considered Category II interfaces.

The Category I jitter generation and transfer criteria are critical to control DS_n jitter accumulation through SONET “islands”. A SONET island is a collection of SONET NEs that create a continuous path for a digital signal that is asynchronously multiplexed at its entry point and asynchronously demultiplexed at its exit point. As a DS_n signal enters and exits these SONET islands, jitter can be mapped from one island to the next by the stuffing

26. “Synchronous DS1 interfaces” refers to DS1 interfaces where an incoming DS1 is byte-synchronously mapped into a VT SPE that is synchronized to the SONET NE’s clock, or where a transmitted DS1 signal is retimed to the NE’s clock (see Sections 3.4.1.1 and 5.4.5.3). Other cases are considered to be Category I.

mechanism in the asynchronous mapping. This can lead to jitter enhancement (i.e., the accumulation of jitter), which may degrade error performance. Islands are an issue for networks as they transition from asynchronous transport to SONET. In an all-SONET network there would be only a single island, and the accumulation of jitter on asynchronous payloads would be less of a concern. In addition, the performance of the synchronization network, as well as the performance of the SONET NEs' internal clocks, may impact the rate at which pointer adjustments are generated and affect jitter accumulation through SONET islands. Section 5.6.2.4 discusses timing jitter for tandem connections of digital equipment in terms of jitter enhancement.

5.6.2.1 Jitter Transfer

Jitter transfer is defined in Section 7 of GR-499-CORE, which contains Category II to Category I, Category I to Category I, and Category II to Category II jitter transfer criteria for non-SONET NEs. In SONET, Category II to Category I jitter transfer (e.g., the OC-N line jitter that appears on an asynchronous DS_n output) is not expected to be significant. Therefore, this section contains only "Category I" criteria (i.e., Category I to Category I jitter transfer, such as that through a multiplexer/demultiplexer pair), and "Category II" criteria (i.e., Category II to Category II jitter transfer, such as that through a SONET regenerator or EDSX²⁷).

Note that measurement methodology (as discussed in GR-499-CORE) is critical for accurate and useful jitter transfer measurements. It is extremely important that the jitter generation characteristics of the NEs and the noise floor of the jitter measurement system be ascertained and taken into account when jitter transfer measurements are made. In all cases, the input jitter amplitudes that are used need to be large enough to result in output jitter amplitudes that are greater than the noise floor of the measurement system and are not significantly affected by generated jitter at the same jitter frequency. If this is not possible at certain jitter frequencies (e.g., if the necessary level of input jitter is greater than the system's jitter tolerance), then any results obtained at those frequencies are erroneous and should be discarded. Note that the noise floor is particularly important at high jitter frequencies, where input jitter levels are limited to lower values by the jitter tolerance mask and the jitter attenuation is required to be large. Jitter transfer tests would normally be expected to concentrate on frequencies within approximately two decades of the break point in the jitter transfer mask.

27. The Category II jitter transfer criteria could also be applied to the OC-N, STS-N electrical, and synchronous DS1 outputs from other line-, loop-, or through-timed SONET NEs. However, if the synchronous outputs are timed from an NE's clocks (which in turn are timed from one or more incoming signals), then the synchronization criteria in Section 5.4 are much more stringent than the jitter transfer criteria in this section, and jitter transfer does not need to be measured. Conversely, if the transmitted signals are directly controlled by the timing of the received signals (as in a SONET regenerator or EDSX) rather than by the NE's clocks, then many of the synchronization criteria are not applicable, and jitter transfer must be measured.

For a system with a linear jitter transfer function, jitter transfer measurements can be made (and identical results can be obtained) using sinusoidal jitter applied to the input signal at any level up to the jitter tolerance level for that interface and that specific jitter frequency. However, SONET systems typically do not have linear jitter transfer functions (both by design and due to inherent factors such as the limited number of stuff opportunity bits available in the asynchronous DS_n to VT or STS SPE mappings), and therefore the results obtained in any jitter transfer tests are likely to depend on the particular input amplitudes used. In general, the primary purpose of the jitter transfer requirements is to prevent performance degradations by limiting the accumulation of jitter through a series of systems such that it does not exceed the network interface jitter requirements (or the jitter tolerance of any of the NEs involved). Thus, it is more important that a system meet the jitter transfer criteria for relatively high input jitter amplitudes (e.g., amplitudes close to the network interface jitter or jitter tolerance limits) than for very low input amplitudes. Therefore, for testing the conformance of a system to the jitter transfer requirements in this document (e.g., to R5-227 [338] or R5-228 [339]), the input jitter amplitude range is limited to 0.1 to 1.0 times the amplitude given by the appropriate jitter tolerance mask. (That is, the jitter transferred through the system must be under the jitter transfer mask for any input jitter amplitude within this range, but is not required to be under the jitter transfer mask for input amplitudes outside of the range.)

5.6.2.1.1 Category I Jitter Transfer

The jitter transfer criteria given in this section limit the amount of jitter on an input DS_n signal at one NE that can be transferred (via the bit stuffing mechanisms used in the asynchronous DS_n mappings) to the output DS_n at another NE. In the development of these criteria, it was assumed that any jitter appearing on an input DS_n signal would be mapped into the STS or VT SPE, and that all of the phase smoothing would occur in the desynchronizer when the DS_n is extracted from the SPE.²⁸ Therefore, the intent of the criteria is to establish the minimum level of phase smoothing that can be performed in desynchronizers. Of course it is also possible for a supplier to design an NE's synchronizer to provide significant attenuation of the input jitter. In a single-supplier network, this could result in acceptable performance even if the NE's desynchronizers do not meet the intent of the criteria. However, for interoperability in a multi-supplier network, a desynchronizer must meet the jitter transfer criteria when it receives a DS_n (within an STS or VT SPE) that was mapped at any synchronizer that conforms to the bit stuffing jitter criteria in Section 3.4.

Note that the criteria in this section are not sufficient to ensure that the desynchronizers in SONET NEs will provide adequate overall jitter attenuation. Specifically, attenuation of

28. Since the bit stuffing mechanism is a form of sampling, input jitter above a certain frequency will be aliased to a lower frequency. For example, in the DS1 mapping, the stuff opportunity bits occur every 0.5 ms, and therefore input jitter at frequencies greater than approximately 1 kHz will be aliased.

the jitter arising from decoded pointer adjustments places more stringent criteria on the desynchronizer characteristics (see Sections 5.6.2.3.2 through 5.6.2.3.5).

- R5-227** [338] For Category I DS_n interfaces other than DS1 and DS3 (e.g., asynchronous DS1C and DS2 interfaces), the jitter transfer function shall meet the jitter transfer requirement in GR-499-CORE, Section 7.3.2.
- R5-228** [339] For Category I DS1 and DS3 interfaces, the jitter transfer function shall be under the mask in Figure 5-26.

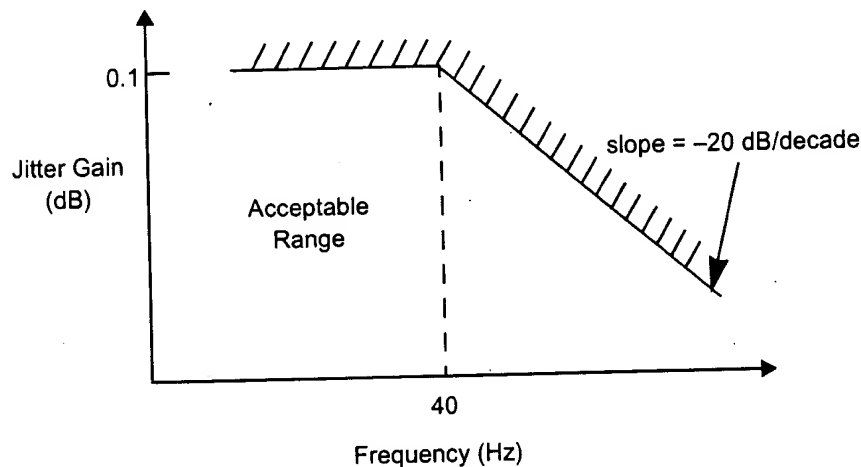


Figure 5-26. Category I DS1 and DS3 Jitter Transfer Mask

Phase smoothing circuits are also needed in cases where a DS_n is demultiplexed from one STS or VT SPE and then multiplexed into another STS or VT SPE within a single SONET NE (e.g., within a DCS that performs DS1 level cross-connections between two STS-1 electrical interfaces), even though the asynchronous DS_n payload does not appear at an external (Category I) interface. These circuits are necessary to filter bit stuffing jitter, intrinsic payload mapping jitter, and pointer adjustment jitter. All of these types of jitter appear at the output of the desynchronizer and, if not filtered, can appear as bit stuffing jitter when the DS_n is multiplexed into a new STS or VT SPE.

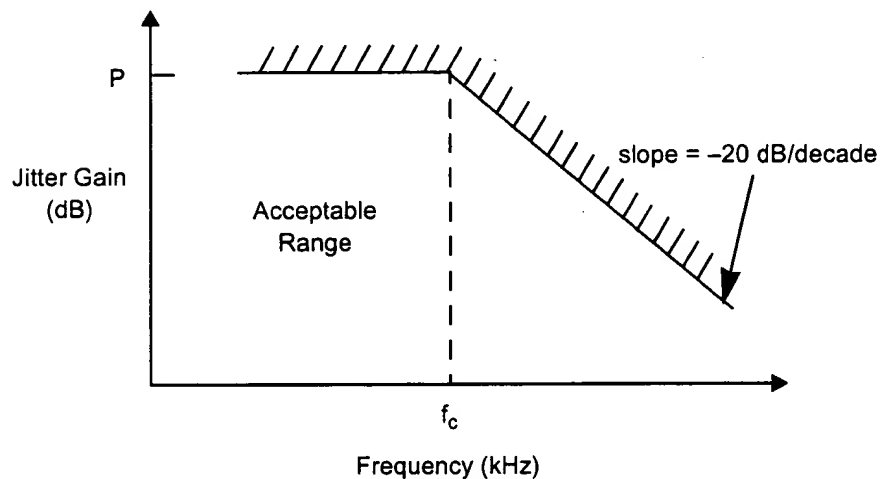
- R5-229** [340] Phase smoothing circuits shall be employed when an asynchronous DS_n payload is demultiplexed from one STS or VT SPE and then multiplexed into another STS or VT SPE within a single SONET NE.

Given the design of some desynchronizers, the most critical type of jitter for an NE to filter is expected to be pointer adjustment jitter. Therefore, more specific criteria on the phase smoothing characteristics of the circuits required by **R5-229 [340]** appear after the pointer adjustment jitter generation criteria, in Section 5.6.2.3.7.

5.6.2.1.2 Category II Jitter Transfer

The jitter transfer requirement in this section limits the amount of jitter on an input OC-N or STS-N electrical signal that can be transferred to the output OC-N or STS-N electrical signal.

R5-230 [341] For Category II interfaces, the jitter transfer function shall be under the mask in Figure 5-27.



OC-N/ STS-N Level	f_c (kHz)	P (dB)
1	40	0.1
3	130	0.1
12	500	0.1
48	2000	0.1

Figure 5-27. Category II Jitter Transfer Mask

5.6.2.2 Jitter Tolerance

5.6.2.2.1 Category I Jitter Tolerance

For Category I interfaces to SONET NEs, the definition of input jitter tolerance given in Section 7.3.1 of GR-499-CORE is the applicable definition.

- R5-231** [342] Category I interfaces to SONET NEs shall meet the Category I input jitter tolerance requirement in Section 7.3.1 of GR-499-CORE.

5.6.2.2.2 Category II Jitter Tolerance

For OC-N interfaces, jitter tolerance is defined as the peak-to-peak amplitude of sinusoidal jitter applied on the input OC-N signal that causes a 1-dB power penalty. This is a stress test intended to ensure that no additional penalty is incurred under operating conditions. For STS-N electrical interfaces, the application of the above definition does not appear to be feasible, so the jitter tolerance definition provided in Section 7 of GR-499-CORE is used for those interfaces.

- R5-232** [343v2] A SONET NE's STS-N electrical and OC-N interfaces, with the exception of OC-48 interfaces that are specified as having "reduced" jitter tolerance (see the discussion below), shall tolerate, as a minimum, input jitter applied according to the mask in Figure 5-28 (with the parameters specified in the figure for the particular rate signal).

SONET NEs with OC-48 interfaces that meet **R5-232** [343v2] are expected to be capable of interworking with all other NEs (in particular, with all SONET regenerators) that meet the OC-48 jitter transfer requirement in Section 5.6.2.1.2. However, in some applications that interworking capability may not be necessary, and the user may only require the NE to meet a reduced jitter tolerance specification (i.e., the NE may not be required to meet **R5-232** [343v2]). The reason for this is that ITU-T Recommendation G.958, *Digital line systems based on the synchronous digital hierarchy for use on optical fiber cables*, has identified a "Type B" regenerator that has tighter jitter transfer characteristics than those required in Section 5.6.2.1.2. If Type B regenerators are used in a SONET line system, then less jitter is transferred through each regenerator. Therefore, the regenerators and the associated terminal interfaces are subjected to less input jitter, and they can operate properly even if they have reduced jitter tolerance.

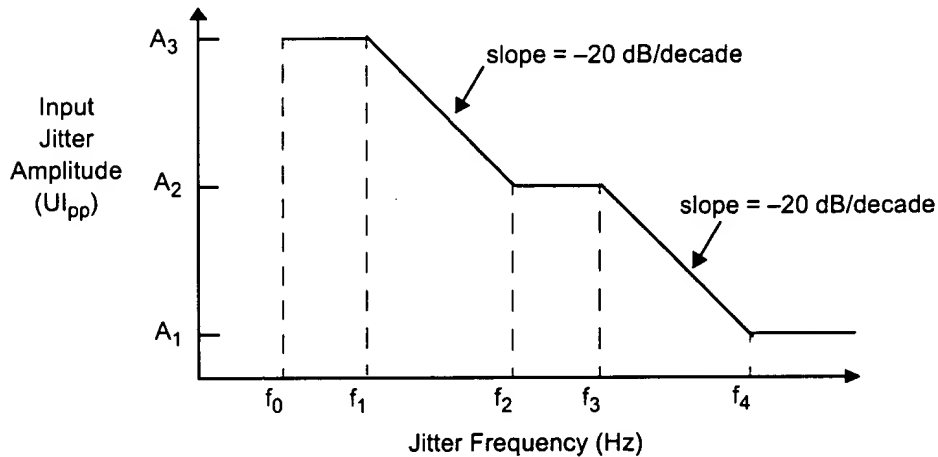
- R5-233** [931] If a SONET NE has reduced OC-48 jitter tolerance, that shall be clearly documented.

Annex A of ANSI T1.105.03 contains the applicable jitter tolerance specification (from ITU-T G.958) for an NE with reduced OC-48 jitter tolerance capabilities. It also contains

the jitter transfer specification (from ITU-T G.958) for Type B regenerators, and discusses how NEs with reduced jitter tolerance may interwork with other equipment subject to certain application constraints.

For Category II DS1 interfaces to SONET NEs, the definition of input jitter tolerance given in Section 7.3.1 of GR-499-CORE is the applicable definition.

- R5-234** [345] Category II DS1 interfaces to SONET NEs shall meet the Category II input jitter tolerance requirement in Section 7.3.1 of GR-499-CORE.



OC-N/STS-N Level	f ₀ (Hz)	f ₁ (Hz)	f ₂ (Hz)	f ₃ (Hz)	f ₄ (Hz)	A ₁ (UI _{pp})	A ₂ (UI _{pp})	A ₃ (UI _{pp})
1	10	30	300	2 k	20 k	0.15	1.5	15
3	10	30	300	6.5 k	65 k	0.15	1.5	15
12	10	30	300	25 k	250 k	0.15	1.5	15
48	10	600	6000	100 k	1000 k	0.15	1.5	15

Figure 5-28. SONET Category II Jitter Tolerance Mask

5.6.2.3 Jitter Generation

Jitter generation is defined in Section 7.3.3 of GR-499-CORE. In SONET, jitter generation criteria exist for both Category I and Category II interfaces. For Category I interfaces the jitter generation criteria are divided into two areas. Section 5.6.2.3.1 contains the criteria in the area of "mapping" jitter generation, while Sections 5.6.2.3.2 through 5.6.2.3.5 contain the criteria in the area of "pointer adjustment" jitter generation. In addition, Section 5.6.2.3.6 contains the Category II jitter generation requirement, and Section 5.6.2.3.7 contains "pointer adjustment jitter to bit stuffing jitter conversion" criteria.

All of the jitter generation tests are performed with no jitter or wander applied at the input. In addition, a bandpass filter is used to limit the jitter generation measurements to the jitter frequency range of interest. For DS_n interfaces (both Category I and Category II), the bandpass filter has a 10-Hz high-pass cutoff frequency with a roll-off of 20 dB/decade, and a low-pass cutoff frequency of at least F_4 (as listed in Figure 7-1 of GR-499-CORE).²⁹ For Category II OC-N and STS-N electrical interfaces, the bandpass filter has a 12-kHz high-pass cutoff frequency with a roll-off of 20 dB/decade, and a low-pass cutoff frequency of at least B_3 (as listed in Table 5-9 for each interface rate).

5.6.2.3.1 Category I Mapping Jitter

Mapping jitter is the sum of the intrinsic payload mapping jitter and the jitter that is generated as a result of the bit stuffing mechanisms used in all of the asynchronous DS_n mappings defined in Section 3.4. Section 3.4 contains bit stuffing jitter criteria (for some of the asynchronous DS_n mappings) that are specifically applicable to the NE that maps the DS_n into a VT or STS SPE, while this section contains criteria that are applicable to a SONET system (e.g., for a signal that is mapped into an SPE at one NE and extracted from that SPE at another NE). To measure the mapping jitter it is necessary for the system to be configured so that no STS or VT pointer adjustments occur during the tests. In a single-product test, this can be accomplished by looping back the OC-N or STS-N electrical signal. In addition, it is necessary to measure the mapping jitter using a variety of DS_n bit rates that meet the DS_n interface criteria in Section 9 of GR-499-CORE.

For some DS_n interfaces (e.g., DS1C and DS2 interfaces), the Category I jitter generation requirement in Section 7.3.3 of GR-499-CORE is currently the applicable requirement. However, tighter requirements have been developed for asynchronous DS1 and DS3 interfaces to SONET NEs, due to concerns about jitter accumulation through SONET islands. In addition, for interoperability in a multi-supplier network, a desynchronizer must meet the mapping jitter generation requirements when it receives a DS_n (within an STS or VT SPE) that was mapped at any synchronizer that conforms to the bit stuffing jitter criteria in Section 3.4.

29. For DS1 interfaces, F_4 is equal to 40 kHz, and for DS3 interfaces it is 400 kHz.

- R5-235** [346] For Category I DS_n interfaces other than DS1 and DS3 interfaces, the mapping jitter shall meet the jitter generation requirement in Section 7.3.3 of GR-499-CORE.
- R5-236** [347] For a Category I DS1 or DS3 interface, the mapping jitter generation shall be less than the value given in Table 5-10.

Table 5-10. Category I Mapping Jitter Limits

Interface	Jitter (UI _{pp})
DS1	0.70
DS3	0.40

5.6.2.3.2 *Category I Jitter Generation Due to Pointer Adjustments, General*

Pointer adjustment jitter criteria are necessary to ensure acceptable DS_n jitter performance during synchronization rearrangements or failure conditions, and for DS_n signals that traverse SONET islands. These criteria are tested using the "pointer test sequences" described by Figures 5-29 through 5-34 in the following sections.

To thoroughly test an NE it is necessary to perform the test sequences a number of times (i.e., using pointer increments, using pointer decrements, using a variety of values for the time variable "T", and using a variety of DS_n bit rates). Table 5-11 gives the limits for T, and along with the notes that follow it, is considered an integral part of the pointer test sequence specifications.

During the pointer adjustment jitter tests, it is important that no buffer spills occur in either the pointer processor circuits or the desynchronizer circuits. The following requirement applies during the single pointer adjustment tests in Section 5.6.2.3.3, the burst pointer adjustment tests in Section 5.6.2.3.4, and the periodic pointer adjustment tests in Section 5.6.2.3.5 that simulate frequency offsets of up to 4.6 ppm in a SONET network. The objective applies during the periodic pointer adjustment tests (Section 5.6.2.3.5) that simulate frequency offset from 4.6 ppm to 20 ppm. In addition, both criteria are applicable during all of the pointer adjustment tests in Section 5.6.2.3.7.

- R5-237** [348] Complete data integrity shall be maintained (i.e., no bit errors shall occur) through the SONET system during all of the pointer adjustment jitter generation tests where T is either in the "required range" given in Table 5-11, or is not applicable (i.e., during the single and burst pointer adjustment tests).

- OS-238'** [349] Complete data integrity should be maintained through the SONET system during all of the periodic pointer test sequences where T is in the "objective range" given in Table 5-11.

Table 5-11. Pointer Test Sequence Parameters

SPE/Payload	t (for burst and periodic tests)	T (Required range for periodic tests)	T (Objective range for periodic tests)
VT1.5/DS1	2 ms	1 s < T < 10 s	0.2 s < T < 1 s
STS-1/DS3	0.5 ms	34 ms < T < 10 s	7.5 ms < T < 34 ms

Pointer Test Sequence Notes:

- Annex C of ANSI T1.105.03 contains measurement methodology information that is useful for these tests.
- Each test sequence consists of an "initialization" period, a "cool down" period, and a "measurement" period. During the initialization period, a series of pointer adjustments (of the same polarity as will be used in the measurement period) is applied to fill or deplete any buffers provided by the NE under test.³⁰ A cool down period of at least 30 seconds is then provided to allow the NE to reach its stable operating condition. Finally, the generated jitter is measured during the measurement period.
- For STS payloads (e.g., DS3), the pointer adjustments are applied to the STS pointers. For VT payloads (e.g., DS1), the pointer adjustments are applied to the VT pointers.
- For all of the sequences, the tests must be performed using a variety of DS_n bit rates that meet the DS_n interface criteria in Section 9 of GR-499-CORE.
- All tests must be run with all positive pointer adjustments, and also with all negative pointer adjustments.
- For periodic sequences, T is constant for each measurement and is determined by the frequency offset between the SPE and its carrier (i.e., the OC-N or STS-N electrical signal in the case of DS3 within an STS-1 SPE). The value of T must be varied over the ranges given in Table 5-11.
- For periodic sequences with added and canceled pointer adjustments, the tests must be run with only added pointer adjustments, and also with only canceled pointer adjustments.

30. Most SONET NEs are expected to include a buffer to absorb many of the pointer adjustments resulting from wander that could occur in the network. For those NEs, pointer adjustment jitter is expected to occur only when the buffer is full (or depleted) and then a pointer decrement (increment) occurs. If the buffer is filled by a series of decrements and then an increment occurs, the buffer would absorb that increment (and also the next decrement since it would no longer be full), and no pointer adjustment jitter would occur.

Although the pointer adjustment patterns described in the following sections are patterns that have been observed in the network and are considered to be relatively likely to occur, a variety of other patterns have also been observed. While some of those patterns (e.g., "dithered" pointer adjustments) have resulted in improved system jitter performance (compared to the patterns described in this document) in both single-product and multi-product configurations, others have resulted in significant performance degradations. Specifically, degradations have occurred when some of the pointer adjustments generated by a pointer processor have been much more closely spaced (in time) than the average spacing.³¹ The following requirements are intended to limit such degradations in two specific steady-state situations. They are not intended to limit the operations that a pointer processor may need to perform to accommodate such things as clock phase transients or wander, network configurations with multiple frequency offsets, or incoming pointer adjustment bursts.

R5-239 **[1025]** In the presence of a constant frequency offset (between the incoming and outgoing signals at a pointer processor) of up to ± 20 ppm and no pointer adjustments in the incoming signal, the minimum spacing (time) between consecutive pointer adjustments generated by a pointer processor shall be greater than or equal to 0.5 times the long-term average spacing.

R5-240 **[1026]** In the presence of incoming pointer adjustments in the patterns shown in Figure 5-32b (VT only), Figure 5-33b (STS only) or Figure 5-34b (STS and VT) and no frequency offset between the incoming and outgoing signals, the minimum spacing between consecutive pointer adjustments generated by a pointer processor shall be greater than or equal to 0.5 times the long-term average spacing.

For the preceding requirements, the time needed to determine the "long-term average" is not specified. However, it generally should be sufficient to monitor the generated pointer adjustments through a complete "cycle" (e.g., 783 STS pointer adjustments, or 104 VT1.5 pointer adjustments), and for frequency offsets and incoming pointer adjustment rates that are large enough that the resulting generated pointer adjustments cannot be considered to be single pointer adjustments (see Section 5.6.2.3.3).

31. For example, in one case a pointer processor that was receiving evenly spaced pointer adjustments on a through STS-1 path was observed to buffer every other incoming adjustment until it received the next adjustment. At that time it would generate two very closely spaced pointer adjustments (e.g., separated by 4 frames), resulting in a pattern of double pointer adjustment bursts at approximately twice the spacing of the original incoming single adjustments.

5.6.2.3.3 Category I Jitter Generation Due to Single Pointer Adjustments

Single pointer adjustments are those that may occur due to wander in the SONET network or small timing reference frequency offsets. For testing purposes it is assumed that the jitter caused by one pointer adjustment will not be affected by the preceding or following pointer adjustments if all of the adjustments are separated by 30 seconds or more.

- R5-241 [350]** The jitter appearing at a Category I DSn interface shall be less than the corresponding value in Table 5-12 when the pointer test sequence in Figure 5-29 is applied.

Table 5-12. Jitter Due to Single Pointer Adjustments

SPE/Payload	Jitter (UI_{pp})
VT1.5/DS1	$A_0 + 0.60^1$
VT2/DS1A	Under Study
VT3/DS1C	Under Study
VT6/DS2	Under Study
STS-1/DS3	$A_0 + 0.30^1$
STS-3c/DS4NA	Under Study

¹ A_0 is the mapping jitter generated by the NE under test

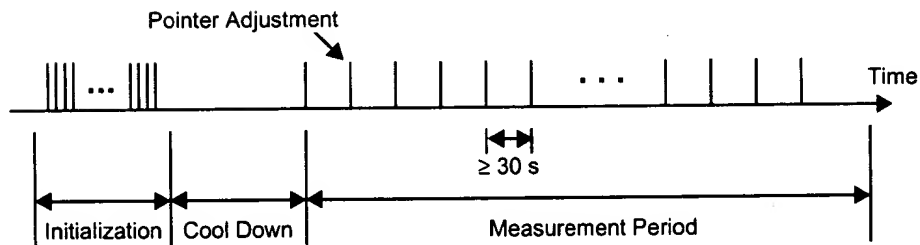


Figure 5-29. Single Pointer Adjustment Test Sequence

5.6.2.3.4 Category I Jitter Generation Due to Bursts of Pointer Adjustments

Jitter generation criteria for bursts of pointer adjustments are currently specified only for DS3 interfaces. Closely spaced pointer adjustments for lower rate signals (e.g., DS1 signals) are not expected to occur due to the larger amount of phase movement necessary to generate VT pointer adjustments. The need for jitter generation criteria for bursts of pointer adjustments for DS4NA payloads is under study.

The pointer adjustment burst sequences specified in the following criteria are intended to simulate the expected worst-case pointer activity due to phase transients caused by synchronization rearrangements (see Section 5.4.4.3.4).

- R5-242** [351] The jitter appearing at a Category I DS3 interface shall be less than $1.3 U_{I_{pp}}$ when the pointer test sequence described in Figure 5-30 is applied.
- R5-243** [352] The jitter appearing at a Category I DS3 interface shall be less than $1.2 U_{I_{pp}}$ when the pointer test sequence described in Figure 5-31 is applied.

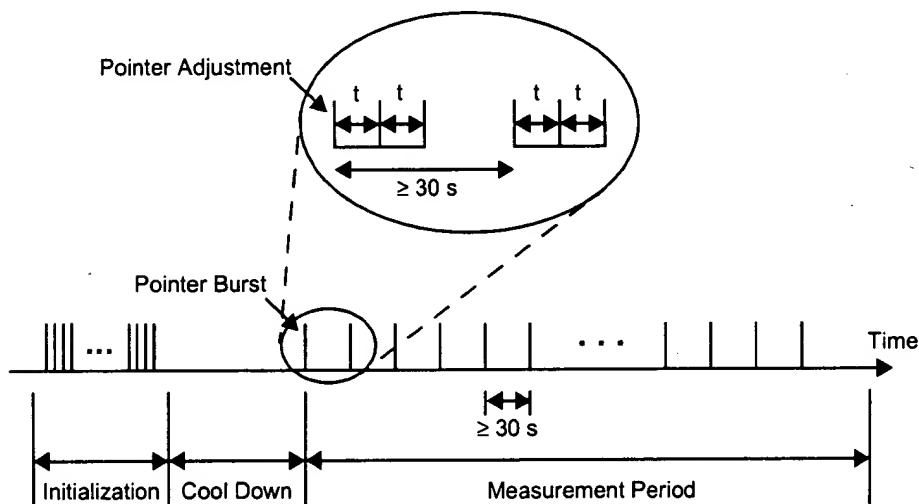


Figure 5-30. Maximum Rate Pointer Burst Test Sequence

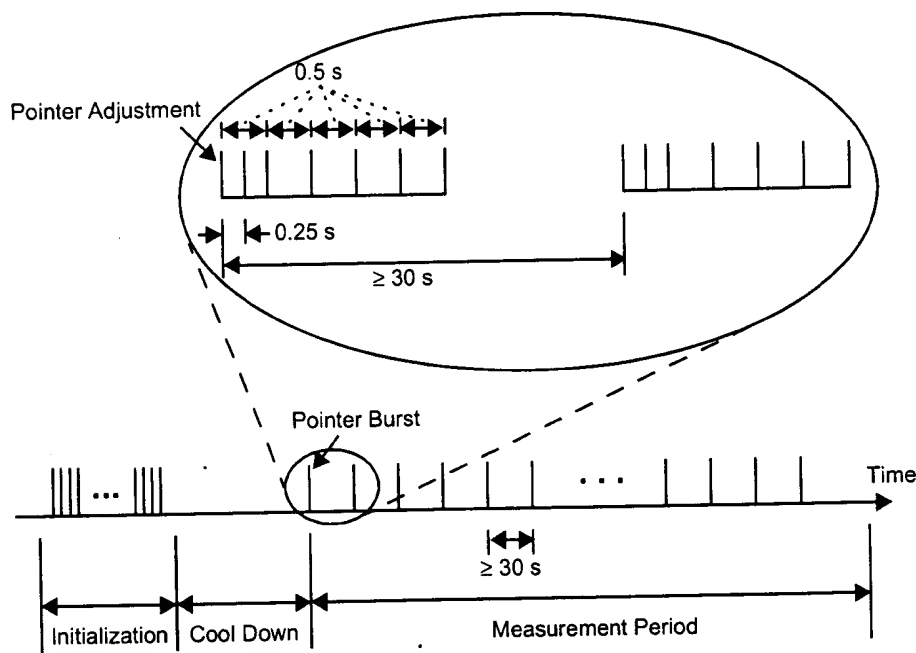


Figure 5-31. Phase Transient Pointer Burst Test Sequence

5.6.2.3.5 Category I Jitter Generation Due to Periodic Pointer Adjustments

The pointer test sequences in this section are intended to simulate the pointer adjustments that could occur in a SONET network in the presence of timing reference offsets. The test sequences specified in the criteria simulate the adjustments that could occur for frequency offsets up to 4.6 ppm (i.e., the “required range” for T shown in Table 5-11), while those in the objectives cover offsets from 4.6 ppm to 20 ppm (i.e., the “objective range”). Currently, criteria exist only for DS1 and DS3 payloads.

R5-244 [353] The jitter appearing at a Category I DS1 or DS3 interface shall be less than the corresponding value given in Table 5-13 when the pointer test sequences described in Figures 5-32, 5-33, and 5-34 are applied with T in the required range.

O5-245 [354] The jitter appearing at a Category I DS1 or DS3 interface should be less than the corresponding value given in Table 5-13 when the pointer test sequences described in Figures 5-32, 5-33, and 5-34 are applied with T in the objective range.

Table 5-13. Jitter Generation Limits for Periodic Pointer Adjustment Sequences

SPE/Payload	Figure Number of Test Sequence					
	5-32b	5-32c and 5-32d	5-33b	5-33c and 5-33d	5-34b	5-34c and 5-34d
VT1.5/DS1	1.3 UI _{pp}	1.9 UI _{pp}	–	–	1.3 UI _{pp}	1.9 UI _{pp}
STS-1/DS3	–	–	1.0 UI _{pp}	1.3 UI _{pp}	1.0 UI _{pp}	1.3 UI _{pp}

It has been observed that some SONET NEs generate pointer adjustments in patterns that do not match any of the pointer test sequences. The following criteria are included (in addition to **R5-239 [1025]** and **R5-240 [1026]**) to ensure acceptable jitter performance for those NEs in single-supplier configurations.

- R5-246** [355] In a single-supplier configuration with a single timing reference offset of 0 to ± 4.6 ppm, the jitter appearing at a Category I DS1 or DS3 interface shall be less than 1.5 UI_{pp}.
- O5-247** [356] In a single-supplier configuration with timing reference offsets equal to twice the specified free-run accuracy of the NEs' internal clocks (e.g., 9.2 ppm for NEs with stratum 3 clocks, 40 ppm for NEs with SMCs), the jitter appearing at a Category I DS1 or DS3 interface should be less than 1.5 UI_{pp}.

5.6.2.3.6 Category II Jitter Generation

- R5-248** [357] The jitter generated at Category II interfaces shall be less than 0.01 UI_{rms}, and shall also be less than 0.10 UI_{pp}.

5.6.2.3.7 Pointer Adjustment Jitter to Bit Stuffing Jitter Conversion

As discussed in Section 5.6.2.1.1, phase smoothing circuits are needed in cases where a single SONET NE demultiplexes a DS_n from one STS or VT SPE and then multiplexes that DS_n into another STS or VT SPE. These circuits are primarily needed to filter pointer adjustment jitter that is generated when the DS_n is demultiplexed, so that it does not appear as bit stuffing jitter when the DS_n is multiplexed into a new STS or VT SPE.

The following criteria are meant to limit the amount of jitter that an NE can convert (internally) from pointer adjustment jitter to bit stuffing jitter, to the same levels as would be applicable if the demultiplexing and multiplexing were performed by separate NEs (i.e., as would be applicable if the DS_n appeared as an external DS_n signal out of one NE, and

then was mapped into a new STS or VT SPE by another NE). The criteria are given in terms of the jitter that would appear at the output of a "nominal" desynchronizer (which is defined as a desynchronizer with filtering characteristics equal to the DS1 and DS3 jitter transfer mask shown in Figure 5-26) when the various pointer test sequences are input to the NE.

- R5-249** [932] The jitter appearing at the output of a "nominal" DS1 or DS3 desynchronizer shall be less than the corresponding value in Table 5-12 when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in Figure 5-29.
- R5-250** [933] The jitter appearing at the output of a "nominal" DS3 desynchronizer shall be less than $1.3 U_{I_{pp}}$ when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in Figure 5-30.
- R5-251** [934] The jitter appearing at the output of a "nominal" DS3 desynchronizer shall be less than $1.2 U_{I_{pp}}$ when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in Figure 5-31.
- R5-252** [935] The jitter appearing at the output of a "nominal" DS1 or DS3 desynchronizer shall be less than the corresponding value given in Table 5-13 when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequences shown in Figures 5-32, 5-33, and 5-34, with T in the required range.
- O5-253** [936] The jitter appearing at the output of a "nominal" DS1 or DS3 desynchronizer should be less than the corresponding value given in Table 5-13 when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequences shown in Figures 5-32, 5-33, and 5-34, with T in the objective range.

5.6.2.4 Jitter Enhancement

Jitter enhancement is defined in Section 7.3.4 of GR-499-CORE, which also contains jitter enhancement criteria for non-SONET NEs. Those criteria help ensure acceptable network performance as jitter is accumulated on a signal that traverses multiple NEs; however, they are not readily translated into criteria for SONET NEs. The possibility of developing SONET jitter enhancement criteria is discussed in GR-253-ILR, Issue ID 253-66.

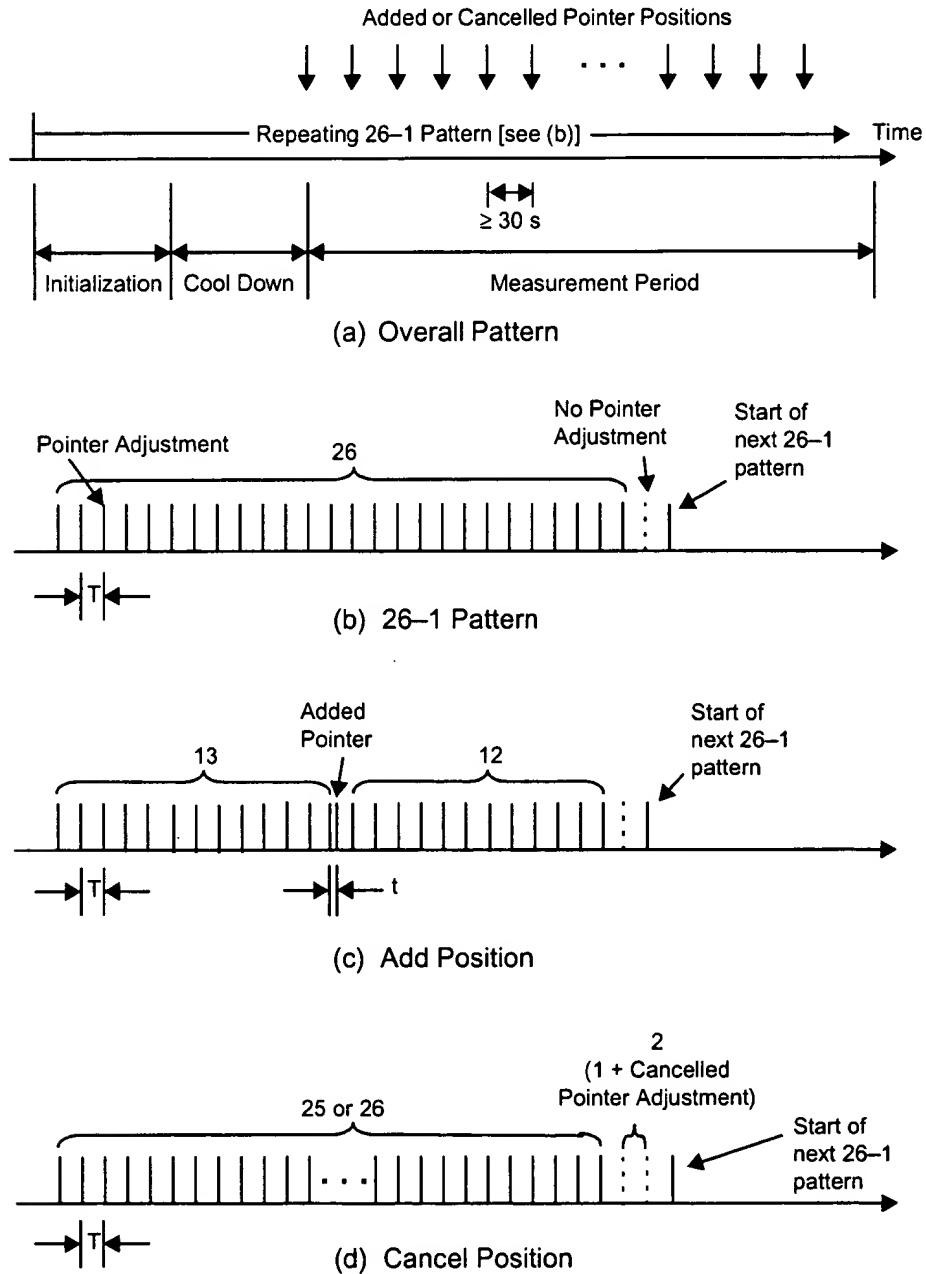


Figure 5-32. Periodic VT1.5 Pointer Adjustment Test Sequence (26-1 Pattern)

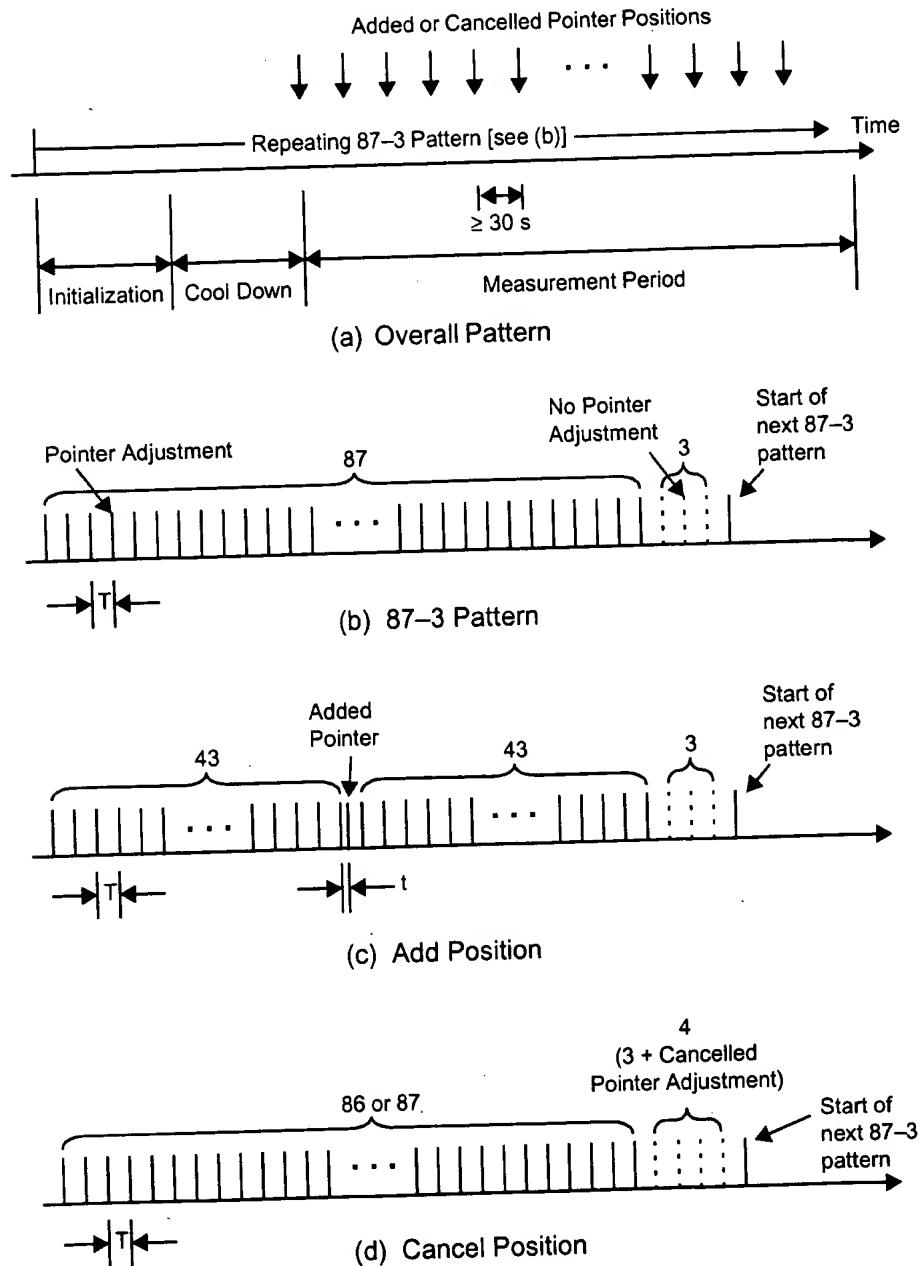


Figure 5-33. Periodic STS Pointer Adjustment Test Sequence (87-3 Pattern)

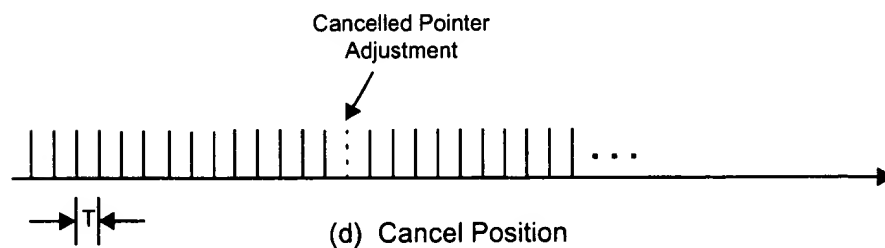
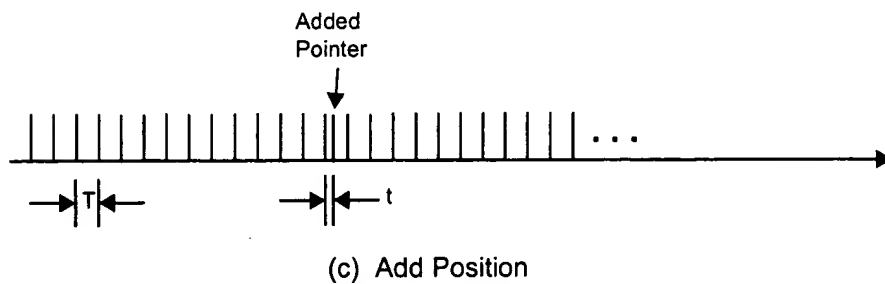
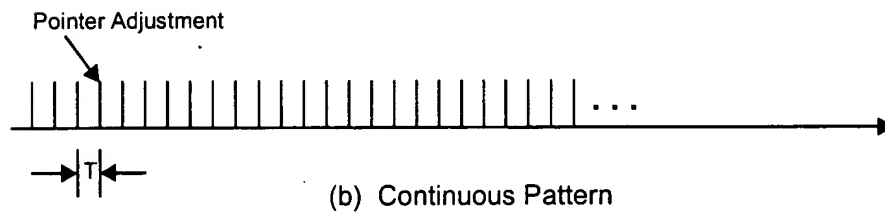
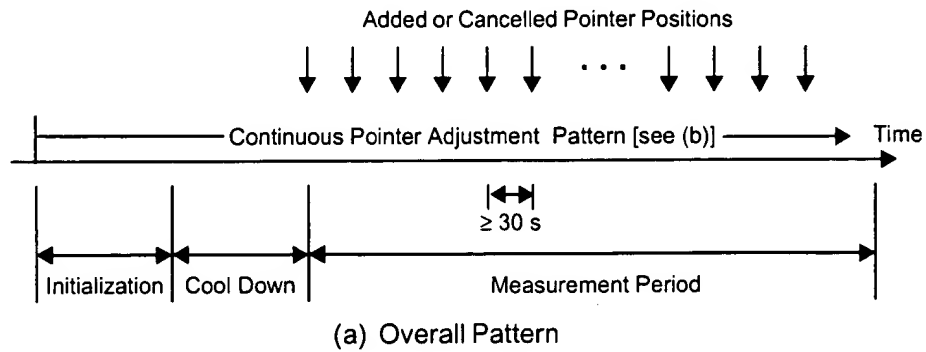


Figure 5-34. Periodic Pointer Adjustment Test Sequence (Continuous Pattern)

5.7 Phase Variations on Payload Signals

This section provides further criteria (in addition to the jitter criteria in Section 5.6) regarding phase variations on DS1 and DS3 payload signals transported on SONET networks. It is based on ANSI T1.105.03a and T1.105.03b (the DS1 and DS3 supplements to ANSI T1.105.03). In addition, those supplements include annexes describing measurement methodology, and those methods are recommended for use in testing SONET NEs against these criteria.

DSn signals carried on SONET networks will incur phase variations due to normal clock noise, any clock offsets (e.g., due to synchronization failures), and the bit-stuffing mechanisms used in the asynchronous DSn to VT or STS SPE mappings. In general, these phase variations must be limited for interworking with existing equipment.

Previously, this section might have been labeled "Wander on Payload Signals." However, the term "wander" is specifically defined to indicate phase variations with frequencies below 10 Hz, and some of the criteria in this section are not bounded by that frequency. Instead, the DS1 payload measurements are made using a 100-Hz first-order low-pass filter, and therefore the more generic term "phase variations" is used. (Note that the DS3 measurements are made using a 10-Hz first-order low-pass filter, and thus could be accurately considered "wander" measurements.)

In all of the phase variation tests, no jitter or wander is applied to the input DS1 or DS3 signal or the timing references to the NEs. In addition, all of the results are given in terms of the MTIE (see Section 5.4.1.2).

5.7.1 Mapping Phase Variations

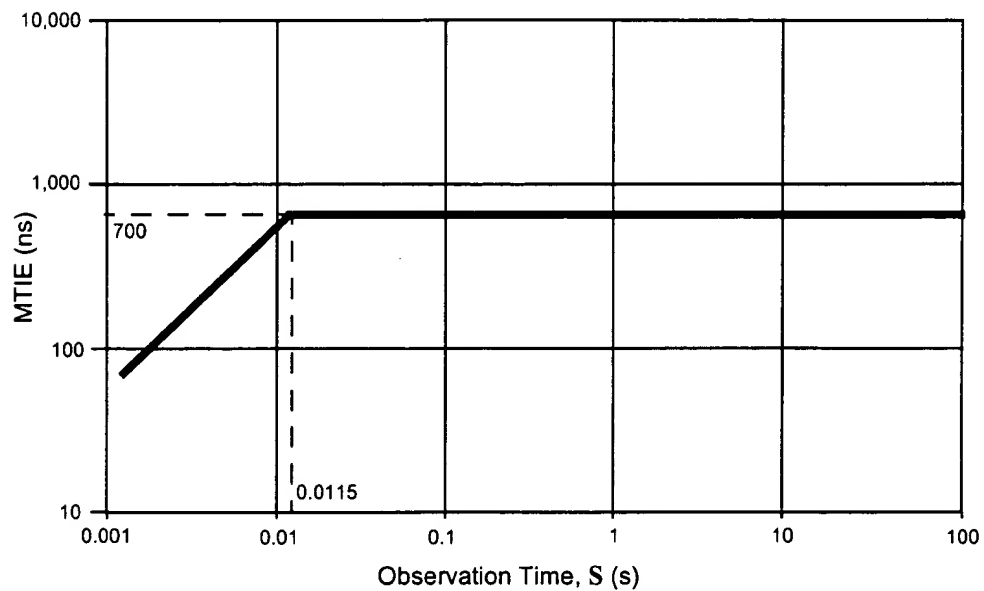
In the asynchronous DS1 to VT1.5 and DS3 to STS-1 mappings, a bit-stuffing mechanism is used to accommodate frequency differences between the DS1 or DS3 and the VT1.5 or STS-1 SPE. This mechanism causes phase variations on DS1 and DS3 signals carried on SONET networks, and these phase variations must be limited for interworking with other network equipment.

To measure the mapping phase variations, it is necessary for the system to be configured so that no STS or VT pointer adjustments occur during the tests. In a single-product test, this can be accomplished by looping back the OC-N or STS-N electrical signal. In addition, it is necessary to measure the mapping phase variations using a variety of DS1 and DS3 bit rates that meet the interface criteria in Section 9 of GR-499-CORE.

R5-254 [358] The mapping phase variations on a DS1 output from a SONET NE shall be below the mask in Figure 5-35.³²

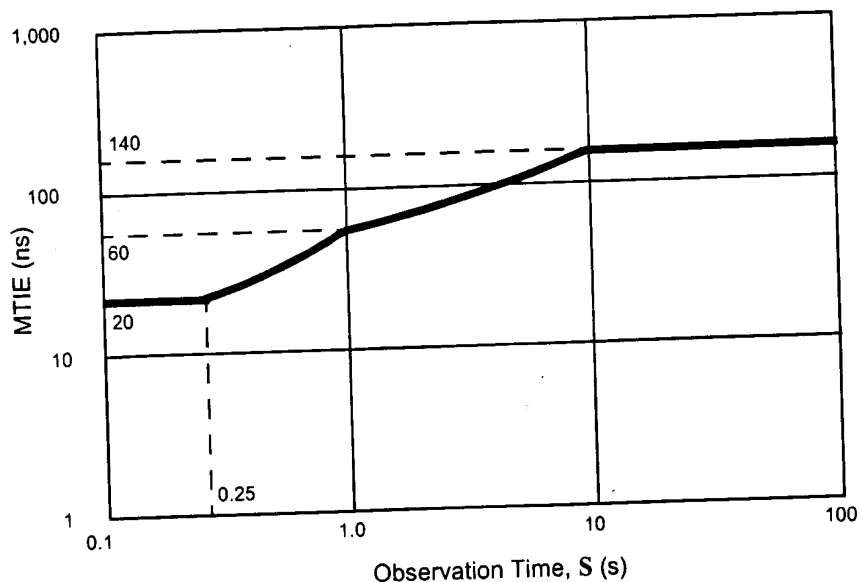
32. The equations for each of the MTIE masks in the figures in this section appear in tables in those figures.

R5-255 [1027] The mapping phase variations on a DS3 output from a SONET NE shall be below the mask in Figure 5-36.



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.001326 < S < 0.0115$	$61000 \times S$
$0.0115 < S$	700

Figure 5-35. DS1 Mapping Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.1 < S$	N/A (jitter region)
$0.1 < S < 0.25$	20
$0.25 < S < 1.0$	$7 + 53 \times S$
$1.0 < S < 10$	$23 + 37 \times S^{0.5}$
$10 < S < 100$	140

Figure 5-36. DS3 Mapping Phase Variation Limits

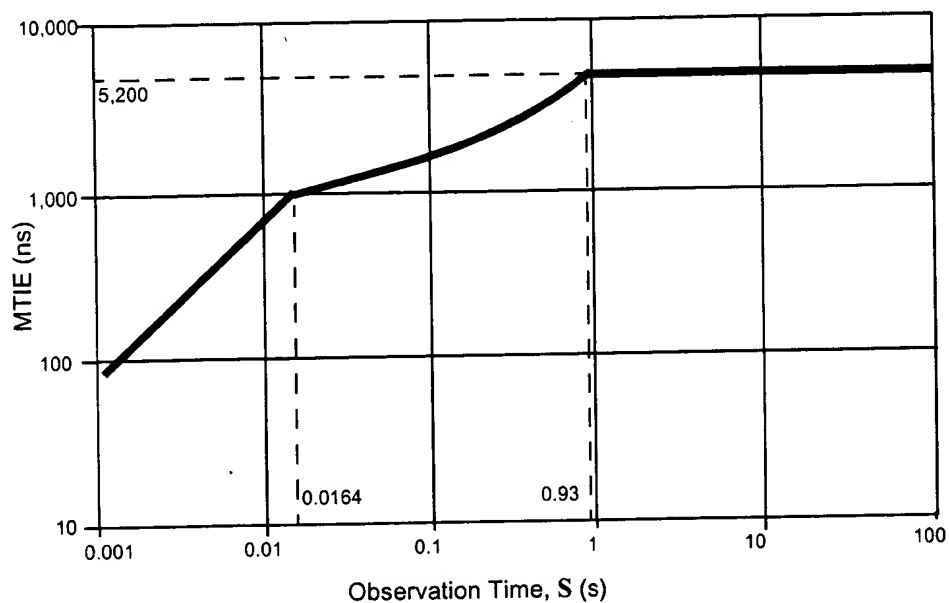
5.7.2 Pointer Adjustment Phase Variations

The pointer adjustment activity in a SONET network is a function of the synchronization characteristics of that network. Clock noise exceeding the buffering capacity of the NEs results in pointer adjustments, which can cause significant phase variations on the DS_n payloads. Since the magnitude of a VT pointer adjustment is much larger than the magnitude of an STS pointer adjustment (e.g., one byte at the VT1.5 rate versus one byte at the STS-1 rate), VT pointer adjustments are expected to occur much less frequently but will cause much larger phase variations. This difference has been verified analytically and via simulation in T1X1.3.

Pointer adjustment statistics can vary widely; however, a set of standard test sequences was developed in T1X1.3 to test the effects of pointer adjustments on the phase variations of DS1 and DS3 payloads. These sequences are the same as the sequences used in the pointer adjustment jitter generation tests in Sections 5.6.2.3.3 and 5.6.2.3.5. As with the jitter tests, the pointer adjustment phase variation tests must be performed with both positive and negative pointer adjustments, and (for the periodic tests) "T" must be varied over the range given in Table 5-11. Unlike the jitter tests, the mapping and pointer adjustment components of the phase variations are expected to be largely separable, and therefore the pointer adjustment tests need to be performed using only a single DS1 or DS3 bit rate (e.g., the bit rate that minimized the mapping phase variations). The results of the mapping phase variation tests (at that bit rate) can then be subtracted from the cumulative results obtained in the pointer adjustment phase variation tests to determine the portion of the MTIE caused specifically by the pointer adjustments.

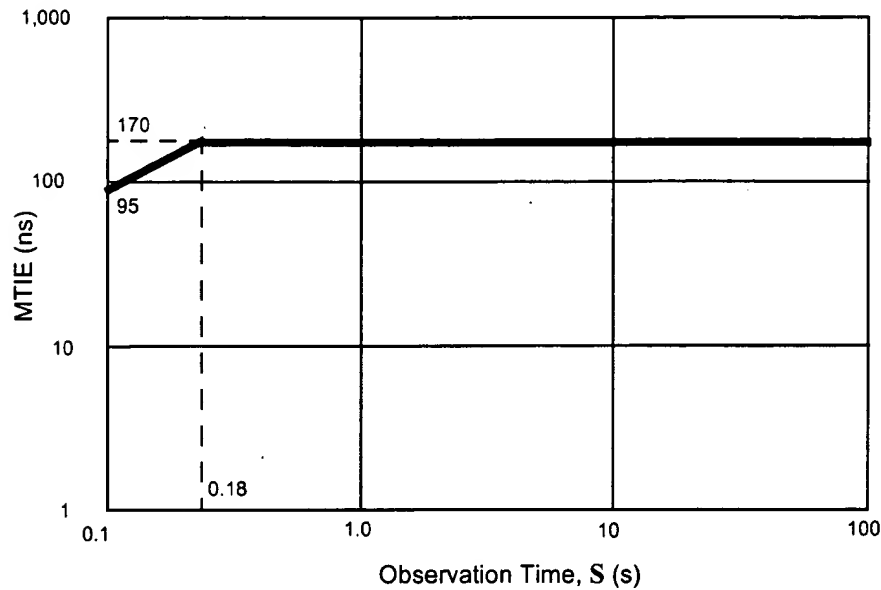
5.7.2.1 Single Pointer Adjustments

- R5-256** [359] The MTIE of a DS1 output from a SONET NE shall be below the mask in Figure 5-37 when the pointer adjustment test sequence in Figure 5-29 is applied.
- R5-257** [1028] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-38 when the pointer adjustment test sequence in Figure 5-29 is applied.



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.001326 < S < 0.0164$	$61000 \times S$
$0.0164 < S < 0.93$	$925 + 4600 \times S$
$0.93 < S$	5200

Figure 5-37. Single VT Pointer Adjustment Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.1 < S$	N/A (jitter region)
$0.1 < S < 0.18$	$945 \times S$
$0.18 < S < 100$	170

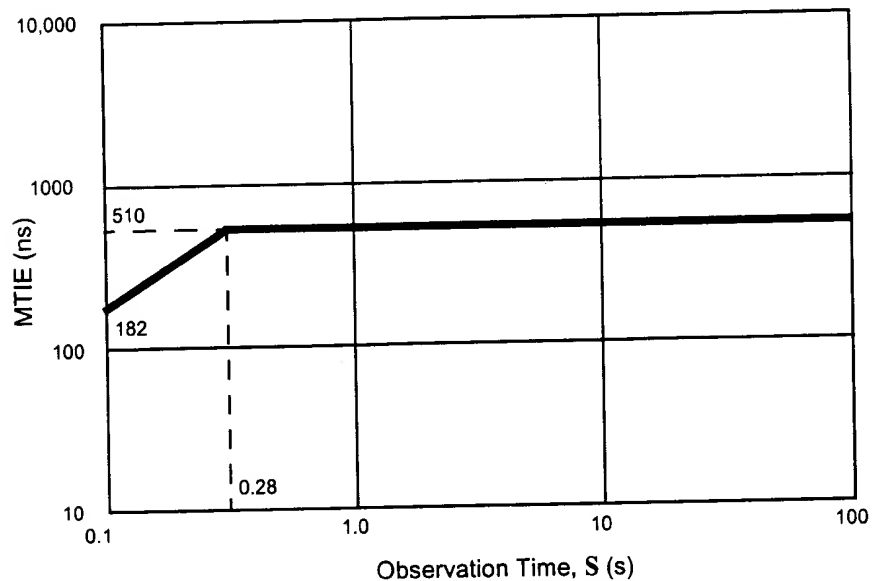
Figure 5-38. Single STS-1 Pointer Adjustment Phase Variation Limits

5.7.2.2 Pointer Adjustment Bursts

As was the case for jitter generation, phase variation criteria for bursts of pointer adjustments are currently specified only for DS3 interfaces. Closely spaced pointer adjustments for lower rate signals are not expected to occur due to the larger amount of phase movement necessary to generate VT pointer adjustments. The pointer adjustment burst sequences specified in the following criteria are intended to simulate the expected worst-case pointer activity due to phase transients caused by synchronization rearrangements (see Section 5.4.4.3.4).

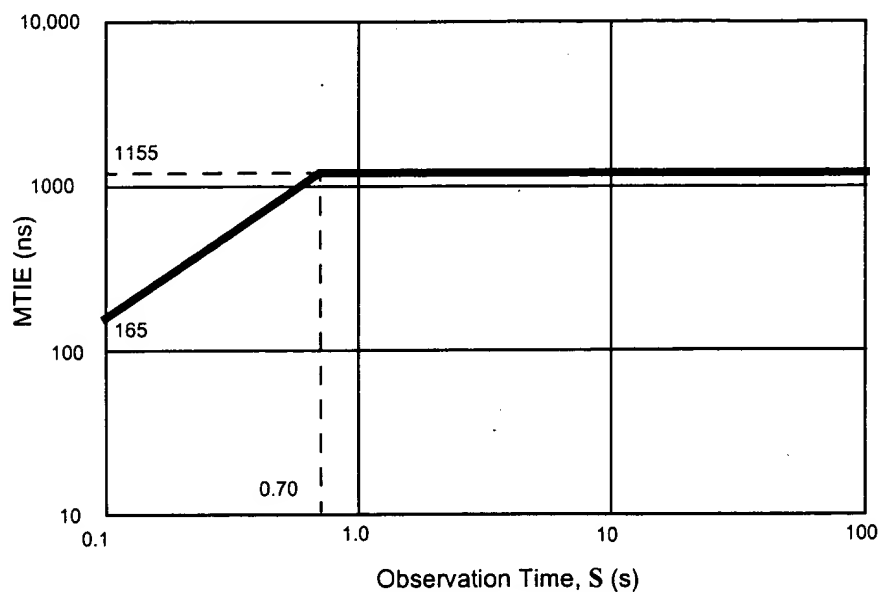
- R5-258 [1029]** The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-39 when the pointer test sequence described in Figure 5-30 is applied.

- R5-259** [1030] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-40 when the pointer test sequence described in Figure 5-31 is applied.



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.1 < S$	N/A (jitter region)
$0.1 < S < 0.28$	$1820 \times S$
$0.28 < S < 100$	510

Figure 5-39. Maximum Rate Pointer Burst Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.1 < S$	N/A (jitter region)
$0.1 < S < 0.70$	$1650 \times S$
$0.70 < S < 100$	1155

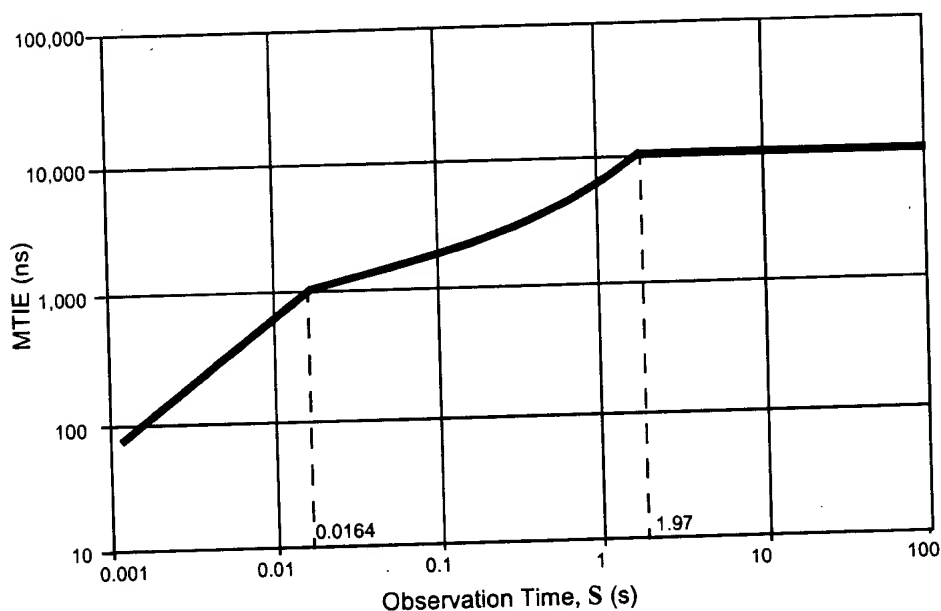
Figure 5-40. Phase Transient Pointer Burst Phase Variation Limits

5.7.2.3 Periodic Pointer Adjustments

Periodic pointer adjustments can be caused by frequency offsets between SONET islands, or by offsets resulting from synchronization reference failures. The magnitude of the offset determines the frequency of the pointer adjustments.

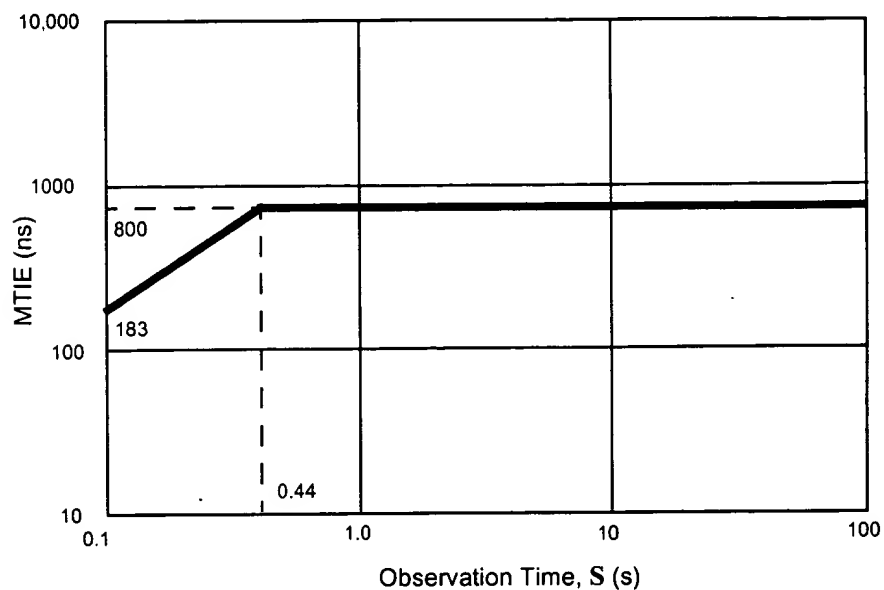
- R5-260** [360v3] The MTIE of a DS1 output from a SONET NE shall be below the mask of Figure 5-41 when the pointer adjustment test sequences in Figures 5-32b and 5-34b are applied with T in the required range (see Table 5-11).

- O5-261** [937v2] The MTIE of a DS1 output from a SONET NE should be below the mask of Figure 5-41 when the pointer adjustment test sequences in Figures 5-32b and 5-34b are applied with T in the objective range.
- R5-262** [1031v2] The MTIE of a DS3 output from a SONET NE shall be below the mask of Figure 5-42 when the pointer adjustment test sequences in Figures 5-33b and 5-34b are applied with T in the required range.
- O5-263** [1032v2] The MTIE of a DS3 output from a SONET NE should be below the mask of Figure 5-42 when the pointer adjustment test sequences in Figures 5-33b and 5-34b are applied with T in the objective range.



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.001326 < S < 0.0164$	$61000 \times S$
$0.0164 < S < 1.97$	$925 + 4600 \times S$
$1.97 < S$	10,000

Figure 5-41. Periodic VT Pointer Adjustment Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$0.1 < S$	N/A (jitter region)
$0.1 < S < 0.44$	$1830 \times S$
$0.44 < S < 100$	800

Figure 5-42. Periodic STS-1 Pointer Adjustment Phase Variation Limits

6. SONET Network Element Operations Criteria

This section contains Operations, Administration, Maintenance, and Provisioning (OAM&P) criteria that are common to SONET NEs. Section 6.1 describes memory administration functions, while Section 6.2 describes maintenance functions including alarm surveillance (Section 6.2.1), Performance Monitoring (Section 6.2.2), and testing (Section 6.2.3). GRs, TRs, and TAs for each type of SONET NE refer to details within this section, and may expand or qualify the description as appropriate for the specific NE.

6.1 Memory Administration

Memory administration deals with the functions necessary to control and administer the databases of NEs. The functions of memory administration discussed in this section include data manipulation, memory backup and restoration, and system administration (including security).

6.1.1 Memory Administration Data

SONET NEs contain provisionable data that is similar to the provisionable data in other NEs, as well as data unique to SONET NEs. For memory administration and other system and network management functions (e.g., fault management), it is necessary to know the termination, cross-connection, and/or multiplex configuration within an NE. Moreover, data specific to a particular signal, such as the mapping of a digital signal payload (asynchronous or byte synchronous), are also necessary.

In the OSI communications environment described in Section 8, an OS and NE can exchange messages by using services and protocols that CMISE provides. The definition of these services depends on an information model that is described in terms of managed object classes and their attributes. An information model for SONET memory administration is an abstraction of the information in a SONET NE that is accessed to perform memory administration functions. GR-836-CORE, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Overview: Transport Configuration and Surveillance for Network Elements* and GR-836-IMD, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Details: Transport Configuration and Surveillance for Network Elements*, define managed object classes, attributes, and associated CMISE service mappings that may be needed to perform memory administration functions. GR-1042-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Overview: Synchronous Optical Network (SONET) Transport Information Model* and GR-1042-IMD, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Details: Synchronous Optical Network (SONET) Transport Information Model*, describe managed object classes and attributes specific to SONET.

For TL1, the NE database can be viewed by the memory administration OS as a collection of logical matrices, or administrative views. Each row of the matrix represents an object entity (i.e., a logical service or resource associated with an NE) of the view, and the columns represent the set of attributes that each object entity may have. Data dictionaries for TL1 object entities are based on administrative views. GR-472-CORE, *OTGR Section 2.1: Network Element Configuration Management* gives a thorough discussion of administrative views for TL1. GR-199-CORE, *OTGR Section 12.2: Operations Application Messages – Memory Administration Messages*, provides administrative views of the data dictionaries for SONET object entities.

6.1.2 Data Manipulation

Data manipulation deals with entering, editing, deleting, and retrieving data in the database of the NE. These functions are required to provision new services and equipment in the network.

- R6-1 [366] Any provisionable feature or parameter shall be provisionable locally by a craftsperson and remotely from an OS.
- R6-2 [361v2] A SONET NE shall use the two-level "Group #, VT #" convention shown in Section 3 of this document (Figures 3-11, 3-13, 3-15, and 3-17) for numbering VTs within an STS-1. This numbering convention shall be used on OS/NE, WS/NE (including any GUI display), and NE/NE interfaces.
- R6-3 [938] A SONET NE shall use either the two-level "STS-3 #, STS-1 #" convention or the single-level "1 to N in order of appearance at the input to the byte-interleaver" convention for numbering STS-1s within an OC-N. These numbering conventions shall be used on OS/NE, WS/NE (including any GUI display), and NE/NE interfaces.

The STS-1 numbering conventions are illustrated in Figures 5-1 and 5-2, while Table 6-1 shows the conversions between the two conventions for OC-3, OC-12 and OC-48 signals. Note that an NE must not use a single-level STS-1 numbering convention where the STS-1s are numbered from 1 to N in the order that they appear in the byte-interleaved OC-N signal [i.e., corresponding to the values in the former C1 (now the J0 and Z0) bytes].

Table 6-1. STS-1 Numbers in OC-N Signals (N = 3, 12, 48)

STS-1 Numbers in an OC-3	
STS-3 #/ STS-1 # Scheme	1 to N Scheme
1,1	1
1,2	2
1,3	3

Order of transmission in a byte-interleaved OC-3 is (1,1), (1,2), (1,3) or 1, 2, 3.

STS-1 Numbers in an OC-12	
STS-3 #/ STS-1 # Scheme	1 to N Scheme
1,1	1
1,2	2
1,3	3
2,1	4
2,2	5
2,3	6
3,1	7
3,2	8
3,3	9
4,1	10
4,2	11
4,3	12

Order of transmission in a byte-interleaved OC-12 is (1,1), (2,1), (3,1), (4,1), (1,2), ..., (4,3) or 1, 4, 7, 10, 2, ..., 12.

STS-1 Numbers in an OC-48			
STS-3 #/ STS-1 # Scheme	1 to N Scheme	STS-3 #/ STS-1 # (cont.)	1 to N (cont.)
1,1	1	9,1	25
1,2	2	9,2	26
1,3	3	9,3	27
2,1	4	10,1	28
2,2	5	10,2	29
2,3	6	10,3	30
3,1	7	11,1	31
3,2	8	11,2	32
3,3	9	11,3	33
4,1	10	12,1	34
4,2	11	12,2	35
4,3	12	12,3	36
5,1	13	13,1	37
5,2	14	13,2	38
5,3	15	13,3	39
6,1	16	14,1	40
6,2	17	14,2	41
6,3	18	14,3	42
7,1	19	15,1	43
7,2	20	15,2	44
7,3	21	15,3	45
8,1	22	16,1	46
8,2	23	16,2	47
8,3	24	16,3	48

Order of transmission in a byte-interleaved OC-48 is (1,1), (2,1), ..., (16,1), (1,2), ..., (16,3) or 1, 4, ..., 46, 2, ..., 48.

- R6-4** [939] For numbering an STS-Mc within an OC-N, a SONET NE shall use the number of the STS-1 in which the STS-Mc starts. This numbering convention shall be used on OS/NE, WS/NE (including any GUI display), and NE/NE interfaces.

This STS-Mc numbering convention is illustrated in Figure 5-2, in which the STS-3c that starts in STS-1 number 4,1 is referred to as "STS-3c Number 4,1" and the STS-12c that starts in STS-1 number 5,1 is referred to as "STS-12c Number 5,1".

6.1.3 Administration of Operations Communications Information

To enable a SONET NE to communicate with management systems and other NEs, the NE needs to be initialized at installation time with communications-related information supplied by the network provider. This information includes options for tailoring each layer of DCC, LCN or OS-NE communications protocol stacks, network address, and target identifiers (TIDs) for NEs employing TL1 messages.

- R6-5** [367] A SONET NE shall provide, via the local craftsperson interface, the ability to initialize the NE with communications-related information upon installation of the NE by the network provider. Such information shall include protocol options for the various operations communications interfaces supported by the NE, the NE's TID (when TL1 messages are supported), and the NE's network address (NSAP) for communications purposes. The supplier shall clearly specify in user documentation those data items that need to be provided upon installation of the NE to ensure proper operations communications involving the NE.

6.1.4 Regenerators

A line with regenerators may contain either physical-layer regenerators or section-terminating regenerators. Regenerators on a line carrying the primary or secondary Section EOC must be section terminating regenerators.

- R6-6** [368] The LTE shall be provisionable to indicate either physical-layer regenerators or section-terminating regenerators are being used for a given line.

Such an indication allows the LTE to identify lines that are suitable for supporting Section EOCs. This is especially useful for LTE with features that allow software reconfigurations of protection systems.

6.1.5 Memory Backup and Restoration

- R6-7** [369] A SONET NE shall provide a local, primary, nonvolatile memory backup.

Such local, primary backup is typically provided within the equipment frame.
GR-472-CORE contains details on memory backup.

A SONET NE may optionally provide an additional (secondary) local, nonvolatile memory backup.

- R6-8** [370] Data shall be backed up in at least one nonvolatile backup memory automatically after each primary memory update.

- R6-9** [371] Restoration of data from the local backup memory, once initiated, shall be completed within 5 minutes.

- O6-10** [372] Restoration of data from the local backup memory, once initiated, should take no more than one minute.

These time requirements represent upper bounds. NE-specific GRs, TRs and TAs specify the time requirements for specific NEs. GR-472-CORE contains further requirements for memory backup and restoration.

There may be catastrophic failure conditions under which even the local nonvolatile memory backup is lost. To guard against such instances, network providers will designate a management application (e.g., a memory administration OS) to be responsible for remotely restoring the memory of NEs so affected. Such management applications must constantly maintain an accurate view of the database changes made in the NE.

- R6-11** [373] A SONET NE shall be able to have its configurable memory restored by a remote memory restoration application identified by the network provider.

- R6-12** [374] A SONET NE shall be able to determine whether or not the source of an update to its configurable memory is the same management application (e.g., OS) as the one responsible for restoring lost primary (and secondary) nonvolatile memory backups.

- R6-13** [375] When the source of a memory update is different from the memory restoration application (e.g., a local craftsperson or other management application), the SONET NE shall send an autonomous indication to the memory restoration application detailing this "hidden update."

- R6-14** [376] If communications to the memory restoration application are not available, "hidden updates" shall be logged by the SONET NE and

reported when asked for by the memory restoration application after communications are restored.

Historically, the remote restoration of an NE's memory has involved the transmission of individual memory administration messages to re-create the last view of the NE's total configuration. While this process can be manageable for small NEs, the process can be quite burdensome and time-consuming for large NEs. Memory restoration via bulk files is becoming a critical operations feature for memory restoration applications and larger NEs.

- CR6-15** [377] A SONET NE may be required to support the ability to have its nonvolatile memory backup restored via bulk file transfer methods by a remote memory restoration management application (e.g., an OS or controller). This feature may be considered an objective or a requirement for certain types of SONET NEs, and suppliers are referred to NE-specific requirements for such instances.

If the SONET NE supports the optional bulk memory restoration feature, the following requirements apply.

- R6-16** [378] If a SONET NE supports the optional bulk memory restoration feature, then the NE shall support the feature via a full 7-layer OSI-based operations interface to a bulk memory restoration application using the memory backup functions and FTAM protocol requirements described in GR-1250-CORE, *Generic Requirements for Synchronous Optical Network (SONET) File Transfer*.
- R6-17** [379] If a SONET NE supports the optional bulk memory restoration feature, then the NE shall be able to report a bulk "snapshot" of its nonvolatile memory backup to the memory restoration application upon request.

GR-836-CORE and GR-836-IMD contain the information model and CMISE service mappings, and GR-199-CORE and GR-833-CORE, *OTGR Section 12.3: Network Maintenance: Network Element and Transport Surveillance*, contain TL1 messages to support transaction-oriented memory backup and restoration. Additional requirements for the management of bulk memory restoration processes are for further study.

6.1.6 System Administration and Security

System administration deals with housekeeping functions needed for proper operation of the NE in the BCC networks. Examples of system administration functions include setting the date and time, and NE identification. System administration functions are provided in GR-472-CORE.

A major area that impacts system administration is security. Security requirements involve routing functions (within the control network), logins, passwords, and security levels (screening options). GR-815-CORE, *Generic Requirements for Network Element/Network System (NE/NS) Security* specifies generic requirements for NE security functions, and TR-NWT-000835, *OTGR Section 12.5: Network Element and Network System Security Administration Messages* specifies TL1 messages that can be used to administer NE security. Preliminary requirements for the information model and service mappings for NE security administration are provided in GR-1253-CORE, *Telecommunications Management Network Security Administration*.

- R6-18** [380] A security mechanism shall be provided within a SONET NE to prevent unauthorized communication to the NE via any ports and communications channels accepting operations-related command inputs,¹ and to allow secure access to the database of the NE. Such a mechanism shall adhere to the security requirements of GR-815-CORE.
- R6-19** [381] The data necessary to support the security mechanism within a SONET NE shall be provided and administered only by authorized security administrators via either WS/NE or OS/NE communications.
- R6-20** [382] A SONET NE shall support security administration functions in conformance with GR-815-CORE.

It is a goal to centralize security administration in a separate security server function, although this capability is not currently fully specified. When it is specified, SONET NEs would be expected to use this security manager function for identification, authentication, and access control.

Additionally, applications in every NE need to employ specific security measures such as logins and passwords to further protect against unauthorized access. Access control features such as authorization levels are needed to limit authorized access to the appropriate functions. The following sections describe further requirements regarding security mechanisms to be provided in SONET NEs.

6.1.6.1 NE Security Mechanism

GR-815-CORE documents the security requirements for NEs and MDs. As GR-815-CORE discusses, a security mechanism implemented within a SONET NE investigates the following questions:

- Is the session requester a valid user?

1. Operations-related commands include functions that allow direct access to any NE resources, such as database, memory, software, processes, etc. In this context, access refers to the following operations: create, retrieve, update, and delete.

- Is the calling address of the request origination in conformity with that of the said valid user?
- Is the user authorized to issue the command being issued?
- Is the user authorized to access the particular portion of the NE database that the command is directed to impact?

Unless all the answers are in the affirmative, a transaction cannot be successfully completed. The security mechanism also provides an audit capability so that unauthorized activities and attempts can be investigated. The mechanism contains the following components:

- Identification and authentication
- System access control
- Resource access control
- Audit.

Identification is the process whereby a session requester's unique and auditable identity (such as the user ID) is recognized. Authentication is the process of verifying the claimed identity of the session requester (such as password check). System access control uses features related to a user's session, such as session "timeout" and real-time detection of password cracking attempts by intruders, thus decreasing the risk of an intruder gaining access by posing as a valid user. Resource access control is enforced after an authorized session has been established, ensuring that no access to the NE database is allowed without proper permission. It serves a dual purpose of accomplishing confidentiality and integrity of the data. Audit features provide the data necessary to support the detection and investigation of unauthorized access.

Bellcore has also issued security requirements for OSI-based TMN interfaces. This document (GR-1469-CORE, *Generic Requirements on Security for OSI-Based Telecommunications Management Network Interfaces*) specifies security requirements for OS-NE, OS-OS, and NE-NE interfaces.

- R6-21** **[383]** A SONET NE supporting interfaces conforming to 7-layer OSI protocol stacks shall conform to the security requirements described in GR-1469-CORE, for layers 1 through 6.
- ...
- When TL1-based interfaces are used in the NE for WS/NE, NE/NE, or OS/NE communications, the NE shall support the data and messages provided in TR-NWT-000835 for the administration of the NE security mechanism.
- ...
- When CMISE interfaces are used on the application layer, the NE shall support data, messages, and mechanisms to conform to the requirements provided by GR-1469-CORE.
-

Work is going on in standards on a security baseline for the interconnection of TMNs. When this work is standardized, TMN administration domains supporting SONET that interconnect may be required to meet this standard.

Data communications is an integral part of the operations and control infrastructure for the SONET architecture. The security of LANs and WANs that interconnect the SONET NEs, OSs, and MDs has a bearing on the security of the SONET transport service. Bellcore has issued security requirements (GR-1332-CORE, *Generic Requirements for Data Communications Network Security*) for the data communications network component of the TMN.

- R6-22** [384] The data communications network component of the SONET TMN shall conform to the security requirements in GR-1332-CORE.

6.1.6.1.1 Identification and Authentication

- R6-23** [385] A SONET NE shall support identification and authentication for all users, for all ports accepting operations-related command inputs, in conformance with GR-815-CORE.

- R6-24** [386] A SONET NE which supports remote access for operations-related command inputs shall provide a feature for additional strong authentication, beyond reusable passwords, such as:

- ... • Third-party authentication
- ... • Public/private key encryption technology

More information on strong authentication within OSI systems can be found in GR-1469-CORE. Additional work is planned to define further details for providing strong authentication in SONET NEs.

- R6-25** [387] When TL1 interfaces are used, a SONET NE shall enforce that a session requester accessing the NE via a TL1/X.25-based OS/NE interface must pass identification information based on X.25 calling address or PVC identifier.

6.1.6.1.2 System Access Control

- R6-26** [388] A SONET NE shall support system access control functions, in conformance with GR-815-CORE.

- R6-27** [389] A SONET NE shall not grant a user remote access unless the user is authenticated via strong authentication.

R6-28 [390] A SONET NE shall employ features corresponding to the timeout interval function that GR-815-CORE describes.

R6-29 [391] A SONET NE shall break down an OSI Application Association if an attempted session request is unsuccessful after a provisionable number of tries. The default number of tries shall be three.

6.1.6.1.3 *Resource Access Control*

R6-30 [392] A SONET NE shall support resource access control and authorization functions in conformance with GR-815-CORE.

R6-31 [393v2] A SONET NE supporting one or more restricted DCCs shall support user identification and access control privileges (see GR-815-CORE) that limit the functionality available to valid outside users accessing the NE via a restricted DCC. The functions allowed via such DCCs shall be definable by local service providers.

6.1.6.1.4 *Audit*

R6-32 [394] A SONET NE shall support audit features, in conformance with GR-815-CORE.

6.1.6.2 *DCC Security*

The location of an NE in the network can necessitate the enforcement of certain security-related restrictions on one or more of the Section DCCs the NE supports. For example, an NE supporting line-side interfaces that cross an administrative boundary needs to apply restrictive message routing functions on the DCCs that cross that boundary. Restrictions are needed to

- Minimize the possibility of unauthorized access to one network's operations functions in NEs and OSs by people, systems, or equipment in another network.
- Limit the broadcast of one network's OS/NE, NE/NE, and WS/NE communications into the other network, thereby reducing the need to control network congestion and also reducing the opportunity to monitor another network's operations communications.

If a SONET network within an administrative domain supports broadcast routing, the network provider may want to disable the DCC crossing the boundary to eliminate excessive broadcasting of messages from outside the routing domain and potential flooding

of the network. However, if a network provider wants an active DCC crossing the boundary to allow, for example, its OS to monitor far-end performance data from the NE in the other domain, the DCC crossing the boundary is labeled restricted and the NE terminating that DCC has to support selective routing. This also eliminates the problem of broadcasting all messages from one domain into another domain. To use the restricted DCC, the boundary NE has to restrict access to the network based on source address and destination address. The network provider should take the particular application and needs into consideration in choosing one of these two methods (i.e., disabled DCC or restricted DCC).

- R6-33** [395] An NE shall be capable of disabling a Section DCC. The default shall be to enable an equipped DCC.
- R6-34** [397] Only properly authorized system administrators shall be allowed to enable or disable an equipped Section DCC.
- R6-35** [398] On reinitialization of the NE after failure, DCCs shall maintain the enabled or disabled state they held before the failure of the NE.
- CR6-36** [399] To provide NE/NE and indirect OS/NE communications paths that may be required by a network provider across administrative boundaries, an NE may be required to terminate one or more enabled DCCs that cross an administrative boundary.
- R6-37** [400] If an NE has the capability to terminate a DCC that crosses an administrative boundary, it shall be capable of classifying it as either restricted or unrestricted for operations security purposes. The default setting shall be unrestricted.
- R6-38** [402] Only properly authorized system administrators shall be allowed to change the classification of a DCC from restricted to unrestricted or vice versa.
- R6-39** [403] A change in classification of a restricted DCC shall not be allowed by communications over that or any other restricted DCC.
- R6-40** [404] On reinitialization of the NE after failure, DCCs shall maintain the restricted or unrestricted states they held before the failure of the NE.

Single-ended maintenance and other functions may still be needed between NEs that span an administrative boundary. An inter-carrier interface (ICI) is an example of a SONET interface that spans an administrative boundary. OSs on one side of the administrative boundary may need, in a tightly controlled and agreed-upon manner, to communicate with an NE on the other side of the boundary via the restricted DCC. Selective routing is one function that must be employed at the network layer in an NE supporting such restricted

DCCs to satisfy security concerns and still allow operation of the links that cross administrative boundaries.

The following requirements apply to a SONET NE having one or more active DCCs that have been classified as restricted:

- R6-41** [405] A SONET NE shall not forward received NPDUs across an administrative boundary to the underlying Data Link Service for a restricted DCC. The NE shall "terminate" the DCC.
- R6-42** [406] A SONET NE shall support a list for each restricted DCC, that itemizes NSAP address pairs that are to be allowed in source and destination address fields of NPDUs allowed into the NE's network.
- R6-43** [407] The list described in **R6-42** [406] shall be established via either WS/NE or OS/NE communications, and only by authorized system administrators via an unrestricted DCC or direct interfaces.
- R6-44** [408] A SONET NE shall screen the NPDUs received from a restricted DCC and only accept the PDU if the source address and destination address matches an allowable source/destination address pair from the list described in **R6-42** [406].
- R6-45** [409] A SONET NE shall enforce access control on the restricted DCC to restrict the requested operations functions that will be permitted on a per-user basis.
- R6-46** [410] A SONET NE shall provide the ability for an authorized administrator to specify the access control privileges assigned to a user for restricted DCC use.

6.1.7 Software Generics

- R6-47** [411] The initial software generic shall be entered in the SONET NE at or before installation.
- O6-48** [412] A SONET NE should be able to receive its initial software load and later software generics via either an OS/NE or NE/NE (or MD/NE) interface using FTAM.

Preliminary communications requirements for software download using the SONET Operations Communications network are provided in GR-1250-CORE. Generic requirements for the support of Remote Software Management are provided in GR-472-CORE.

R6-49 [413] The NE shall provide the ability to retrieve (locally via a WS and remotely via an OS) the current version ID of software.

O6-50 [414] Software updates, or patches, should be identified and included in the current version report.

6.1.8 Self-Inventory

Self inventory (also referred to as "automatic discovery") is a feature of an NE to inventory its own equipment.²

O6-51 [415] A SONET NE should be able to report to a management application or a craftsperson its equipage (including plug-ins, common equipment, and software), option settings, and its crossconnect configuration.

6.2 Maintenance

This section provides SONET NE maintenance criteria necessary to maintain the NE and the network. Maintenance requirements include alarm surveillance, performance monitoring (PM), testing, and control features that are essential in the normal operation of the NE. The requirements in this section address functions that are used to perform the following maintenance tasks:

- **Trouble detection** deals with detecting defects and declaring failures. Section 6.2.1.1 defines defects and failures related to SONET.
- **Trouble or repair verification** is the process of verifying the continued existence or nonexistence of a problem before beginning or closing out work on that problem.
- **Trouble sectionalization** deals with sectionalizing the failure to one of the terminating NEs or the facility that connects them. This process is achieved through alarms, maintenance signals (e.g., AIS), PM data, test access, and loopbacks.
- **Trouble isolation** deals with the isolation of failures down to a replaceable circuit pack, module, or a fiber. Test access, loopbacks, performance data, and diagnostics available within the SONET NEs are used to achieve this isolation.
- **Restoration** permits service to be restored even though the failure may not have been repaired. Protection switching, described in Section 5.3 and rerouting of traffic are examples of how restoration can be achieved.

2. Detailed NE requirements for self-inventory functions are for further study.

In formulating generic maintenance criteria common to all SONET NEs, it is useful to take a functional view of equipment. Thus, maintenance criteria common to all SONET NEs may be specified in terms of STE, LTE, STS PTE, and VT PTE.

Suppliers are referred to SR-NWT-002723, *Applicable TL1 Messages for SONET Network Elements*, for information on how specific TL1 messages apply to the maintenance functions discussed here.

6.2.1 Alarm Surveillance

Alarm surveillance deals with the detection and reporting of certain degraded conditions in the network. This section enumerates those conditions that shall be detected within the SONET signal and the NE. In addition to defining the various conditions, this section also describes the NEs' actions in response to detecting those conditions.

This GR refers to occurrences in the network that are detected as "defects". A defect is defined to be a limited interruption in the ability of an item to perform a required function, and SONET NEs are required to "detect" and "terminate" certain defects on the incoming signals relevant to the layers of functionality they provide. Detection of a defect may cause a particular action to be performed (e.g., the transmission of a maintenance signal), while termination of the defect generally causes that action to be halted (e.g., removal of the maintenance signal). When a defect persists for a period of time (i.e., a soaking interval), a corresponding failure is generally "declared" and the NE sets a failure indication. Once a failure indication has been set, if the defect is terminated and remains absent for a period of time, then the failure is "cleared". Failure indications may or may not be automatically reported to the OS, and the reported indications may be alarmed or nonalarmed. Failure indications, whether or not they are automatically reported to the OS, are available and retrievable by the OS or some other user system interface. Some failure indications may also result in audible and/or visible alarm indications locally at the NE. For more details on the failure monitoring process and alarm strategy, including definitions of critical, major, and minor alarms, see GR-474-CORE, *OTGR Section 4: Network Maintenance: Alarm and Control for Network Elements*, and GR-820-CORE.

The defects and failures discussed in Section 6.2.1.1 are considered to be directly detected defects and failures and represent root-cause problems on the incoming SONET bitstream. Sections 6.2.1.2, 6.2.1.3, and 6.2.1.4 describe symptomatic defects, which are detected via maintenance signals on the incoming SONET bitstream that are generated as a result of an upstream or downstream SONET NE detecting one of the directly detected defects described in Section 6.2.1.1. A general model of defect detection and failure declaration is illustrated in Figure 6-1.

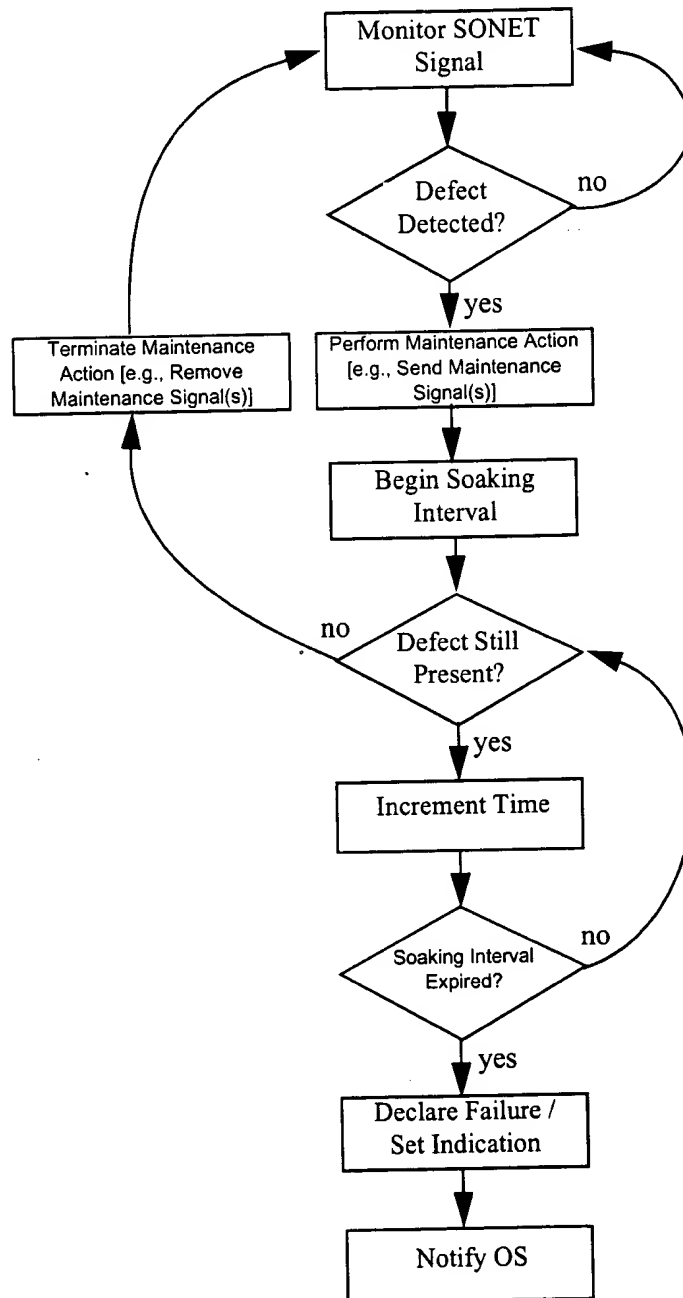


Figure 6-1. General Defect Detection and Failure Declaration Model

6.2.1.1 Directly Detected Defects and Failures

6.2.1.1.1 Loss of Signal (LOS)

To detect a failure that occurs at the source (e.g., laser failure) or the transmission facility (e.g., fiber cut), all incoming SONET signals are monitored for loss of physical-layer signal (optical or electrical). The detection of an LOS defect must take place within a reasonably short period of time for timely restoration of the transported payloads.

- R6-52** [416v2] A SONET NE shall monitor all incoming SONET signals (before descrambling) for an "all-zeros patterns," where an all-zeros pattern corresponds to no light pulses for OC-N optical interfaces and no voltage transitions for STS-1 and STS-3 electrical interfaces. An LOS defect shall be detected when an all-zeros pattern on the incoming SONET signal lasts 100 μ s or longer. If an all-zeros pattern lasts 2.3 μ s or less, an LOS defect shall not be detected.

The treatment of all-zeros patterns lasting between 2.3 μ s and 100 μ s for the purpose of LOS defect detection is not specified and is therefore left to the choice of the equipment designer. For testing conformance to the LOS detection requirement, it is sufficient to apply an all-zeros pattern lasting at most 2.3 μ s, and to apply an all-zeros pattern lasting at least 100 μ s.

In addition to monitoring for all-zeros patterns a SONET NE may also detect an LOS defect if the received signal level (e.g., the incoming optical power) drops below an implementation-determined threshold.

- O6-53** [940] If a SONET NE monitors the received signal level for the purpose of detecting LOS defects, then its signal level threshold should be selected such that an LOS defect will not be detected if the BER is still acceptable (e.g., no LOS defect if the BER is better than the SF BER threshold used for protection switching in linear APS, see Section 5.3.3.1).
- R6-54** [417] The SONET NE shall terminate an LOS defect when the incoming signal has two consecutive valid framing alignment patterns and, during the intervening time (one frame), no all-zeros pattern qualifying as an LOS defect exists.
- R6-55** [418v2] A SONET NE shall declare an LOS failure when the LOS defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an LOS failure indication and send an alarm message to an OS.
- R6-56** [419] For the purposes of trunk conditioning, a SONET NE that contains DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels
-

in byte-synchronously mapped DS1s shall declare an LOS failure if it is subject to a period of short, intermittent LOS defects. For failures resulting from the NE intermittently detecting and terminating the defect, the NE shall integrate the time during which the defect persists, using a 4:1 to 15:1 count-up/count-down ratio. During a sustained LOS defect, the integrator shall count up to reach the alarm threshold in 2.5 (± 0.5) seconds. Upon reaching the alarm threshold, the NE shall declare the LOS failure, set an LOS failure indication, and send an alarm message to an OS. If the defect is terminated before the threshold is reached, the integrator shall count down at a slope 1/4 to 1/15 of the count-up slope.

SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s are also required to use this integration technique for declaring certain other failure (e.g., see R6-64 [427] for LOF failures). In addition, SONET NEs may also apply this (or some similar) integration technique to declare failures due to intermittent defects in applications where trunk conditioning is not required.

R6-57 [420] A SONET NE shall clear an LOS failure when the LOS defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the SONET NE shall clear the LOS failure indication and send a clear message to an OS.

R6-58 [421v2] SONET NEs interfacing with DS1, DS1C, DS2, or DS3 signals shall detect LOS on those signals according to the requirements in GR-499-CORE.

Criteria for LOS on other transport signal interfaces is for further study.

6.2.1.1.2 Loss of Frame (LOF)

R6-59 [422] All incoming SONET signals shall be monitored for LOF. A SONET NE shall detect an LOF defect when an SEF defect (see Section 5.5) on the incoming SONET signal persists for 3 ms.

A SONET NE may optionally implement a 3-ms integration timer to deal with intermittent SEFs when monitoring for LOF. Such a 3-ms integration timer consists of an SEF timer and an in-frame timer that operate as follows:

1. The in-frame timer is activated (accumulates) when an SEF defect is absent. It stops accumulating and is reset to zero when an SEF defect is detected.
2. The SEF timer is activated (accumulates) when an SEF defect is present. It stops accumulating when the SEF defect is terminated. It is reset to zero when the SEF defect is absent continuously for 3 ms (i.e., the in-frame timer reaches 3 ms).

3. An LOF defect is detected when the accumulated SEF timer reaches the 3-ms threshold. Once detected, the LOF defect is terminated when the in-frame timer reaches 3 ms.

- R6-60** [423] If the optional integration timer described above is provided for LOF monitoring, the supplier shall clearly describe its use in the product documentation.
- R6-61** [424] The SONET NE shall terminate an LOF defect 1 ms to 3 ms after terminating the SEF defect on the incoming SONET signal, if the SEF defect is not (re)detected before the LOF defect is terminated.
- O6-62** [425] The SONET NE should terminate an LOF defect 1 ms after terminating the SEF defect on the incoming SONET signal, if the SEF defect is not detected within the 1-ms time period. (This objective is not applicable if the optional 3-ms integration timer, described above, is used.)
- R6-63** [426v2] A SONET NE shall declare an LOF failure when an LOF defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOF failure, the NE shall set an LOF failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies (see Section 6.2.1.8).
- R6-64** [427] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare LOF failures. Upon declaring an LOF failure, the NE shall perform the actions listed in **R6-63 [426v2]**.
- R6-65** [428v2] A SONET NE shall clear an LOF failure when the LOF defect is absent for 10 (± 0.5) seconds. Upon clearing the LOF failure, the SONET NE shall clear the LOF failure indication and send a clear message to an OS (if the failure was reported to the OS).

For SONET signals, detection of the LOF defect initiates certain maintenance-related actions (e.g., generation of AIS and RDI), which are discussed in Sections 6.2.1.2, 6.2.1.3, and 6.2.1.4. For DS_n signals, it is the detection of an OOF that initiates those maintenance-related actions. Requirements on DS_n OOF detection, in-frame detection, and maintenance-related actions for DS1, DS1C, DS2, and DS3 signals are contained in GR-499-CORE. Framing criteria for other transport signal interfaces are for further study.

- R6-66** [429] An NE shall monitor for DS_n OOF on DS_n paths that are terminated by the NE.

While a SONET NE may terminate the DS_n path in some applications, in many other applications it is expected to provide clear-channel transport of DS_n payloads.³ In clear-channel DS_n transport, incoming DS_n signals are monitored for DS_n LOS defects, incoming SONET signals are monitored for a variety of defects (e.g., LOS, LOF), and the appropriate AIS is inserted downstream when any of those defects are detected. In some cases an NE providing clear-channel DS_n transport may also non-intrusively monitor the framing of the DS_n signals (e.g., to support the accumulation of intermediate DS_n path PM, see Section 6.2.2.9); however, in those cases AIS is not inserted when a DS_n OOF is detected. If an NE inserts AIS when it detects a DS_n OOF, then it is not providing clear-channel transport. As indicated in Section 3.4, the DS_n asynchronous mappings were defined specifically for the clear-channel transport of DS_n signals in STS and VT SPEs. Therefore, in most situations where the asynchronous mappings are used, a SONET NE is not expected to insert AIS upon detecting DS_n OOF.

- R6-67** [941] If an NE supports an asynchronous DS_n mapping as defined in Section 3.4, then it shall be capable of providing clear-channel transport of DS_n signals using that mapping.

In addition, an NE that supports an asynchronous DS_n mapping could support one or more options to provide non-clear-channel transport of DS_n signals using that mapping. In non-clear-channel transport, the NE monitors for DS_n OOF (as if it were terminating the DS_n path) and inserts the appropriate AIS downstream when an OOF is detected. If such a feature is provided, the criteria in GR-499-CORE on detecting DS_n OOF and declaring DS_n OOF failures are applicable.

- CR6-68** [942] An NE that supports an asynchronous DS_n mapping may be required to provide non-clear-channel transport where the incoming DS_n signal is monitored for DS_n OOF.

- CR6-69** [943] An NE that supports an asynchronous DS_n mapping may be required to provide non-clear-channel transport where the outgoing DS_n signal (i.e., the DS_n signal that is demultiplexed from the STS or VT SPE) is monitored for DS_n OOF.

A SONET NE may also provide an option to monitor any error detection bits provided in the DS_n signal (e.g., the P-bits in a DS3 signal) and correct those bits when errors are detected. That option, and the support of non-clear-channel DS_n transport in general, may be useful in situations where the DS_n signals are subsequently transported on certain older transmission systems.⁴ Note however, that clear-channel DS_n transport is the preferred mode for NEs supporting the asynchronous DS_n mappings, and that **CR6-68 [942]** and

3. Note that in this context, the term "clear-channel transport" refers to the transport of DS_n payloads through the SONET network. The same term is also used in other contexts to refer to the use of various line codes or consecutive-zero suppression codes which allow DS_n signals with arbitrary numbers of consecutive logical zeros to be transported on DS_n facilities (e.g., B8ZS).

CR6-69 [943] are not expected to be changed to requirements in future issues of this document.

6.2.1.1.3 Loss of Pointer (LOP)

To detect a failure related to the pointer processing mechanism (e.g., for trouble isolation purposes), an NE that processes pointers also monitors incoming STSs and VTs for LOP. SONET equipment detects an LOP defect when a valid STS or VT pointer cannot be obtained by using the pointer interpretation rules described in Section 3.5 and the pointer word does not contain all-ones. An LOP defect is also detected when a number of consecutive pointers with the NDF set to '1001' but not indicating concatenation is received. [Under normal operation, the NDF would be set only once to indicate a change in pointer value (or VT size). Except when set continuously in a concatenation indicator, consecutive NDFs would indicate a pointer processor failure (e.g., stuck bits).]

R6-70 [430v3] STS PTE and LTE that processes the STS pointers shall monitor for LOP-P. An LOP-P defect shall be detected if a valid pointer is not found in N consecutive frames (where $8 \leq N \leq 10$), or if N consecutive NDFs (other than in a concatenation indicator, see Section 3.5.1.3) are detected. An LOP-P defect shall not be detected when LTE is receiving and relaying an all-ones STS pointer, or when STS PTE is receiving pointers that qualify as those necessary to cause the detection of an AIS-P defect (i.e., three or more consecutive all-ones pointers).

R6-71 [944v2] VT PTE and STS PTE that processes VT pointers shall monitor for LOP-V. An LOP-V defect shall be detected if a valid pointer is not found in N consecutive superframes (where $8 \leq N \leq 10$), or if N consecutive NDFs are detected. An LOP-V defect shall not be detected when STS PTE is receiving and relaying an all-ones VT pointer, or when VT PTE is receiving pointers that qualify as those necessary to cause the detection of an AIS-V defect.

Several definitions of "valid pointer" are possible, and the particular definition that a supplier uses will have a large impact on how that supplier interprets the above requirement and implements an LOP detection algorithm. However, since LOP is not expected to be a common defect in the network, these different definitions and interpretations are not expected to have a significant impact on network performance.

The intended definition of "valid pointer" is:

4. For example, some older asynchronous fiber optic transport systems required incoming DS3 signals to have valid framing and (in some cases) parity. In those cases, the SONET NE would need to insert DS3 AIS when it detects a DS3 OOF, and may also need to correct the P-bits in the DS3 signal.

1. A pointer with an in-range value, the N-bits set to their normal value, and, for VT pointers, constant or correct size bits (see Section 3.5.2.3) that is received identically in three consecutive frames (superframes)
2. A pointer containing the concatenation indicator (if the NE is able to accommodate concatenated signals, see Section 3.5.1.4) that is received identically in three consecutive frames.

Therefore, the intended interpretation of the LOP defect detection requirement is that a pointer processor detects an LOP defect at the end of any N frame (superframe) window ($8 \leq N \leq 10$) that does not contain at least one set of three consecutive frames (superframes) with either a valid pointer as defined above, or all-ones pointer words (in which case the all-ones pointer would be relayed or an AIS defect would be detected instead).

Other possible definitions of a "valid" pointer could include a single pointer that exactly matches the current "valid" pointer value, or a pointer that causes the receiving NE to locate the start of the STS or VT SPE at a new location according to the pointer interpretation rules in Sections 3.5.1.6 and 3.5.2.6. For example, a received pointer that is interpreted as an increment or a decrement from the previous "valid" pointer value could establish a new "valid" pointer value.

Although various definitions and interpretations of the requirement are acceptable, some limitations are needed. Therefore, as a minimum, a pointer processor must detect an LOP defect if it receives ten consecutive frames (superframes) in which all of the following are true:

1. None of the pointers contain the same pointer value as the current "valid" pointer value, along with the N-bits set to their normal value, and for VT pointers, constant or correct size bits
2. None of the pointers can be interpreted as indicating an increment or decrement from the current "valid" pointer value (thus establishing a new current valid pointer value)
3. No three consecutive frames (superframes) contain a constant, in-range pointer value, the N-bits set to their normal value, and for VT pointers, constant or correct size bits
4. None of the pointers have the NDF set (along with an in-range pointer value, and, for VT pointers, constant or correct size bits), or all of the pointers have the NDF set
5. For LTE that processes STS pointers (STS PTE that processes VT pointers) and meets the all-ones pointer relay objective (see Sections 3.5.1.5 and 3.5.2.5), none of the pointers is all-ones
6. For STS (VT) PTE, or LTE that processes STS pointers (STS PTE that processes VT pointers) and does not meet the all-ones pointer relay objective, no three consecutive frames (superframes) contain all-ones pointers.

- R6-72** [431v2] STS PTE and LTE that processes the STS pointers shall terminate an LOP-P defect when the STS has a valid pointer with a normal NDF, or a valid concatenation indicator, in three consecutive frames.
- R6-73** [432v2] STS PTE shall terminate an LOP-P defect when it detects an AIS-P defect.
- R6-74** [945] LTE that processes STS pointers shall terminate an LOP-P defect when it relays an all-ones STS pointer.
- R6-75** [946] VT PTE and STS PTE that processes VT pointers shall terminate an LOP-V defect when the VT has a valid pointer with a normal NDF in three consecutive superframes.
- R6-76** [947] VT PTE shall terminate an LOP-V defect when it detects an AIS-V defect.
- R6-77** [948] STS PTE that processes VT pointers shall terminate an LOP-V defect when it relays an all-ones VT pointer.
- R6-78** [433v2] A SONET NE shall declare an LOP-P failure when an LOP-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOP-P failure, the NE shall set an LOP-P failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
- R6-79** [949] A SONET NE shall declare an LOP-V failure when an LOP-V defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOP-V failure, the NE shall set an LOP-V failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
- R6-80** [434] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare LOP-P and LOP-V failures. Upon declaring an LOP-P or LOP-V failure, the NE shall perform the actions listed in **R6-78 [433v2]** or **R6-79 [949]**.
- R6-81** [435v2] A SONET NE shall clear an LOP-P failure when an LOP-P defect is absent for 10.0 (± 0.5) seconds. Upon clearing the LOP-P failure, the SONET NE shall clear the LOP-P failure indication and send a clear message to an OS (if the failure was reported to the OS).
- R6-82** [950] A SONET NE shall clear an LOP-V failure when an LOP-V defect is absent for 10.0 (± 0.5) seconds. Upon clearing the LOP-V failure, the
-

SONET NE shall clear the LOP-V failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.1.4 Equipment Failures

This GR does not define equipment failure states, since they are largely implementation-dependent. However, it does list a minimum set of conditions to be reported as alarms.

- R6-83** [436] Equipment failures shall be classified as either Service-Affecting (SA) or Non-Service-Affecting (NSA), depending on whether they affect the services that the equipment transports.
- R6-84** [437] Equipment failures shall be classified as critical, major, or minor.
- R6-85** [438] Because hardware designs vary, the report of the equipment failure shall describe the failure condition.
- R6-86** [439] The NE shall be able to declare the following equipment failures (as a minimum):
- ... • Fuse or power circuit failures
 - ... • Synchronization equipment failures
 - ... • Protection switching equipment failures
 - ... • CPU failures
 - ... • Local nonvolatile backup memory failures
 - ... • SONET signal origination and termination equipment failures
 - ... • Receiver failures (optical detector failures)
 - ... • Transmitter failures (light source failure, including laser failures)
 - ... • Non-SONET signal (e.g., DSn) origination and termination equipment failures
 - ... • Switching matrix module failures (if cross-connect functionality is provided)
 - ... • DCC hardware failures (also see Section 6.2.1.1.7)
 - ... • Manual removal of in-service (i.e., active) equipment.
- R6-87** [440] Upon declaring an equipment failure, a SONET NE shall

- ... • Switch to duplex or standby equipment, if available
- ... • Set a local indication
- ... • Send an alarm message to an OS.

In addition, certain equipment failures cause AIS to be generated (see Section 6.2.1.2), and an NE may declare and report equipment failures that are not specifically listed in this section.

R6-88 [441] Upon clearing an equipment failure, a SONET NE shall clear the equipment failure indication and send a clear message to the OS.

CR6-89 [442] A SONET NE may be required to detect and report certain environmental conditions in some applications (e.g., NEs in a CEV).

The CPU of a SONET NE may be provided in a redundant manner so that a standby can be automatically switched into service upon failure of the main CPU. NE-specific GRs, TRs, and TAs contain requirements regarding this hardware redundancy feature for the CPU and other equipment (e.g., DCC or LCN termination).

O6-90 [443] A SONET NE should detect operating system or other software errors, and report them to an OS, independently of CPU hardware failures.

6.2.1.1.5 *Loss of Synchronization*

Synchronization equipment failures are listed with other equipment failures in Section 6.2.1.1.4, which covers the declaration and clearing of all equipment failures. This section describes the Loss of Synchronization failure, which may or may not be caused by a synchronization equipment failure.

R6-91 [444] A SONET NE shall have the capability of declaring a Loss of Synchronization failure, due to either loss of primary timing reference or loss of secondary timing reference. Upon declaring the failure, the NE shall set a Loss of Synchronization failure indication and send a message to an OS. The message shall include an indication of reference switches and the reason for failure (e.g., LOS, LOF or OOF, synchronization message, etc.).

R6-92 [445] Upon clearing the Loss of Synchronization failure, a SONET NE shall clear the Loss of Synchronization failure indication and send a clear message to the appropriate OS.

Section 5.4 discusses criteria for loss of synchronization, including the use of synchronization status messages. Section 8 of GR-1244-CORE describes additional operations requirements for synchronized clocks.

6.2.1.1.6 APS Troubles

Four types of failures related to the operation of the APS channel are defined for LTE that supports linear APS. These are Protection Switching Byte failures (Section 6.2.1.1.6.A), Channel Mismatch failures (Section 6.2.1.1.6.B), APS Mode Mismatch failures (Section 6.2.1.1.6.C), and Far-end Protection Line failures (Section 6.2.1.1.6.D). In addition, Section 6.2.1.1.6.E contains a requirement concerning the alarm generated by an NE when it receives an AIS-L and cannot perform a protection switch.

The criteria in Sections 6.2.1.1.6.A, 6.2.1.1.6.B, and 6.2.1.1.6.D apply to LTE that is *operating* in a linear APS mode other than the 1+1 unidirectional mode, while the criteria in Section 6.2.1.1.6.C apply to LTE that is *provisioned to operate* in a linear APS mode other than the 1+1 unidirectional mode. For example, LTE that is provisioned to operate in the 1:1 bidirectional mode, but that is actually operating in the 1+1 unidirectional mode (because that is the mode indicated by the far-end LTE), must meet the criteria in Section 6.2.1.1.6.C, but does not need to meet those in Sections 6.2.1.1.6.A, 6.2.1.1.6.B, and 6.2.1.1.6.D.

Since all LTE that uses linear APS must transmit the appropriate codes in the APS channel, LTE that is operating in the 1+1 unidirectional mode can expect to receive those codes, and may (but is not required to) use them to detect and declare the defects and failures discussed in the following sections. Note however, that those defects and failures must not affect the operation of the 1+1 unidirectional system, and that the LTE must meet **R5-87 [207]** and **O5-88 [208]** (even though it is not operating in the bidirectional mode) to avoid declaring extraneous alarms in some situations.

6.2.1.1.6.A Protection Switching Byte Failure

Unless it is operating in the 1+1 unidirectional mode, LTE must monitor the incoming K1 byte for Protection Switching Byte failures. A Protection Switching Byte defect occurs when either an inconsistent APS byte or an invalid code is detected. An inconsistent APS byte occurs when no three consecutive K1 bytes of the last 12 successive frames are identical, starting with the last frame containing a previously consistent byte. An invalid code occurs when the incoming K1 byte contains an unused code (see Table 5-4), or a code irrelevant for the specific switching operation (e.g., Reverse Request while no switching request is outstanding) in three consecutive frames. An invalid code also occurs when the incoming K1 byte contains an invalid channel number in three consecutive frames.

R6-93 **[446]** LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Protection Switching Byte defect within

50 ms of the occurrence of either an inconsistent APS byte or an invalid code, unless the condition for terminating the defect occurs.

- R6-94** [447] LTE shall not detect a Protection Switching Byte defect when it has detected an AIS-L defect.
- R6-95** [448] LTE shall terminate the Protection Switching Byte defect within 50 ms of the occurrence of three consecutive, identical, and valid APS codes, unless the condition for detecting the defect occurs.
- R6-96** [449] LTE shall not terminate a Protection Switching Byte defect when it has detected the AIS-L defect.
- R6-97** [450] An NE shall declare a Protection Switching Byte failure when a Protection Switching Byte defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall perform the following actions:
- ... 1. A Protection Switching Byte failure indication shall be set and an alarm message shall be sent to an OS.
 - ... 2. If a working channel that was being selected from the protection line is switched back to the working line (see Section 5.3.5.5), a message shall be sent to an OS indicating the switch back to the working line.
- R6-98** [451] An NE shall clear the Protection Switching Byte failure when the Protection Switching Byte defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Protection Switching Byte failure indication and send an alarm clear message to an OS.

6.2.1.1.6.B *Channel Mismatch Failure*

Unless it is operating in the 1+1 unidirectional mode, LTE must monitor the channel numbers in its transmitted K1 byte (bits 5 through 8) and the received K2 byte (bits 1 through 4) for Channel Mismatch. (In the 1+1 unidirectional mode, each end operates independently and the LTE does not need to monitor for Channel Mismatch.) If the channel number in the received K2 byte is not identical to the channel number transmitted in the K1 byte, there is a mismatch.

Under normal conditions, a mismatch will occur each time the LTE changes the channel number on its transmitted K1 byte (i.e., when the channel with the highest priority switch request changes). A mismatch could also occur if the channel number on the incoming K2 byte changes (e.g., due to a failure at the far-end LTE).

- R6-99** [452] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Channel Mismatch defect if the channel

numbers in the transmitted K1 byte and the received K2 byte do not match for 50 ms.

- O6-100** [453] LTE operating in a linear APS mode other than the 1+1 unidirectional mode should detect a Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte are mismatched in three consecutive frames.

Section 5.3.5.4 discusses the actions that LTE is required to take when it detects or terminates a Channel Mismatch defect.

- R6-101** [454] LTE shall not detect a Channel Mismatch defect when it has detected an AIS-L defect.
- R6-102** [455] LTE shall terminate the Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte match in three consecutive frames.
- R6-103** [456] LTE shall not terminate a Channel Mismatch defect when it has detected an AIS-L defect.
- R6-104** [457] An NE shall declare a Channel Mismatch failure if the Channel Mismatch defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a Channel Mismatch failure indication and send an alarm message to an OS.
- R6-105** [458] An NE shall clear the Channel Mismatch failure if the Channel Mismatch defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Channel Mismatch failure indication and send an alarm clear message to the OS.

6.2.1.1.6.C APS Mode Mismatch Failure

Unless it is provisioned to operate in the 1+1 unidirectional mode, LTE must monitor the mode of operation indicators in the incoming K2 byte for APS Mode Mismatch failures. As discussed in Section 5.3.5.2.3, the APS mode information in bits 5 through 8 of the received K2 byte is mismatched when either:

1. LTE provisioned for 1+1 protection switching receives an indication from the far-end LTE (in bit 5 of K2) that it is provisioned for 1:n (or vice versa), or
2. LTE provisioned for bidirectional switching receives an indication from the far-end LTE (in bits 6-8 of K2) that it is provisioned for unidirectional switching (or vice versa).

Codes other than '101' and '100' in the incoming K2 bits 6-8 are ignored for the purposes of monitoring for the APS Mode Mismatch defect.

- R6-106** [459] LTE provisioned to operate in a linear APS mode other than the 1+1 unidirectional mode shall detect an APS Mode Mismatch defect within 100 ms of receiving the first of five consecutive samples of frames (which may or may not be consecutive frames) with identical mode information (either in bit 5 of K2 or bits 6-8 of K2) that is mismatched, as defined above, unless the condition for terminating the defect occurs before the defect is detected.
- R6-107** [460] LTE shall not detect an APS Mode Mismatch defect when it has detected an AIS-L defect.
- R6-108** [461] LTE shall terminate an APS Mode Mismatch defect within 50 ms of receiving the first of five consecutive samples of frames (which may or may not be consecutive frames) with identical mode information that is not mismatched as defined above, unless the condition for detecting the defect occurs before terminating the defect.
- R6-109** [462] LTE shall not terminate an APS Mode Mismatch defect when it has detected an AIS-L defect.
- R6-110** [463] An NE shall declare an APS Mode Mismatch failure when an APS Mode Mismatch defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set the APS Mode Mismatch failure indication and send an alarm message to an OS.
- R6-111** [464] An NE shall clear the APS Mode Mismatch failure if the APS Mode Mismatch defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the APS Mode Mismatch failure indication and send an alarm clear message to the OS.

6.2.1.1.6.D *Far-End Protection-Line Failure*

Unless it is operating in the 1+1 unidirectional mode, LTE must monitor the K1 byte for Far-End Protection-Line failures. When LTE receives an SF code for the protection line, it knows that the far-end LTE is no longer receiving its request, or (in bidirectional operation) that the far-end LTE considers the near-end LTE's request to be invalid. In unidirectional operation, the near-end LTE can maintain any existing switch (e.g., continue to select a working or extra traffic channel from the protection line) if the far-end LTE maintains and indicates the appropriate bridge in K2 bits 1 through 4 (see **R5-77 [198]**). In bidirectional operation, any working channel that had been switched to the protection line

is switched back to its working line (since the SF on protection request is higher priority than any other request except Lockout of Protection).

- R6-112** [465] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Far-End Protection-Line defect when it receives three consecutive K1 bytes with the code indicating "SF on the protection line."
- R6-113** [466] LTE shall terminate the Far-End Protection-Line defect when it receives three consecutive, identical, and valid K1 bytes with any code other than "SF on the protection line."
- R6-114** [467v2] An NE shall declare a Far-End Protection-Line failure when a Far-end Protection-Line defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, the NE shall set a Far-End Protection-Line failure indication and (if it is a Reported failure) send an alarm message to an OS.
- R6-115** [468] An NE shall clear the Far-End Protection-Line failure if the Far-End Protection-Line defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Far-End Protection-Line failure indication and send an alarm clear message to the OS (if the failure was reported to the OS).

6.2.1.1.6.E Other APS Failures

Although they are not defined as defects/failures (see Section 6.2.1), BER-based SF and SD conditions that are used for protection switching purposes can have a direct impact on the transport of traffic, and therefore the user needs to be alerted when those conditions exist. The following requirements are intended to support this. Note that to avoid the generation of multiple autonomous messages in situations where a BER-based SF or SD condition is repetitively detected and cleared over a relatively short time period, the following requirements include "soaking times" similar to those used in most of the failure declaration and clearing criteria in the various subsections of Section 6.2.1.

- R6-116** [1064] When a BER-based SF condition persists for 2.5 (± 0.5) seconds, an NE shall set an SF BER indication and send an alarm message to an OS.
- R6-117** [1065] When a BER-based SD condition persists for 2.5 (± 0.5) seconds, an NE shall set an SD BER indication and send an alarm message to an OS.
- R6-118** [1066] An NE shall clear an SF BER indication 10 (± 0.5) seconds after detecting that the BER is less than the SF clearing threshold, assuming that it does not detect that the BER is greater than the SF threshold during that time period (see Section 5.3.4). Upon clearing an SF BER indication, the

NE shall send an alarm clear message to the OS (if the SF BER was reported to the OS).

- R6-119** [1067] An NE shall clear an SD BER indication 10 (± 0.5) seconds after detecting that the BER is less than the SD clearing threshold, assuming that it does not detect that the BER is greater than the SD threshold during that time period (see Section 5.3.4). Upon clearing an SD BER indication, the NE shall send an alarm clear message to the OS (if the SD BER was reported to the OS).
- R6-120** [1068] As a default, the alarm level for SF BER and SD BER indications shall be Not Alarmed.
- R6-121** [469] If LTE that provides the linear APS feature is unable to perform automatic protection switching upon receiving an AIS-L signal, it shall send an SA alarm to an OS, indicating the inability to perform a protection switch.

6.2.1.1.7 DCC Failure

DCC hardware failures are listed with other equipment failures in Section 6.2.1.1.4, which covers the declaration and clearing of all equipment failures. This section describes the DCC failure in more detail.

A DCC failure is defined as either a DCC hardware failure (as in Section 6.2.1.1.4) or failure of the line carrying the DCC. In the first case, standby DCC equipment can be provided to protect against a service DCC equipment failure. In the second case, the DCC protection scheme described in Section 8.3.1.3 is used to recover from a failure on the line carrying the primary DCC. ES-IS and IS-IS routing protocols can be used as a protection scheme when alternate routes are available to route messages around failures where both primary and secondary DCCs either have been lost or are unavailable.

- R6-122** [470] A SONET NE shall be able to declare a DCC failure. Upon declaring the failure, it shall set a DCC failure indication and send a message to an OS.
- R6-123** [471] Upon clearing a DCC failure, a SONET NE shall clear the DCC failure indication and send a clear message to an OS.

6.2.1.1.8 *Signal Label Mismatch*

A received STS or VT signal label (the C2 byte or V5 bits 5 through 7, respectively) is considered mismatched if it does not equal either a label value corresponding to the locally provisioned PTE functionality or the label value corresponding to the equipped, non-specific code (see Tables 3-2, 3-3, and 3-4). Two types of defects are currently defined at each path layer. These are the Payload Label Mismatch (PLM) and Unequipped (UNEQ) defects. Tables 6-2 and 6-3 identify the conditions corresponding to these two types of defects. In those tables, the "Received Payload Label" corresponds to the signal label in the STS and VT paths received on the incoming signal, while the "Provisioned Functionality" corresponds to the mapping used by the STS or VT PTE. Only in-service, provisioned PTE can detect these defects. PTE is considered provisioned if it has been configured for a mapping (or only supports one mapping), and has been assigned a timeslot (or is hardwired to a specific timeslot) in a SONET signal. When an UNEQ or PLM defect is detected, the appropriate AIS is sent to downstream equipment and an ERDI code (if supported) is sent to upstream equipment.

Note that if STS PTE has detected an AIS-P or LOP-P defect on its incoming signal, the C2 byte cannot be accessed to monitor for PLM-P or UNEQ-P defects. Therefore, it can neither detect nor terminate a PLM-P or UNEQ-P defect. Similarly, if VT PTE has detected an AIS-V or LOP-V defect, the V5 byte cannot be monitored for PLM-V or UNEQ-V defects (see Section 6.2.1.8.2).

6.2.1.1.8.A *STS Payload Label Mismatch*

- R6-124** [472] STS PTE shall detect an STS Payload Label Mismatch (PLM-P) defect within 250 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of mismatched STS Signal Labels (C2 byte), as specified in Table 6-2.

For testing conformance to this requirement, it is sufficient to transmit a continuous string of identical, mismatched STS Signal Labels in the C2 byte for at least 250 ms.

- O6-125** [473] STS PTE should detect a PLM-P defect immediately upon receipt of five contiguous frames with mismatched STS Signal Labels, as specified in Table 6-2.

- R6-126** [476] STS PTE shall terminate a PLM-P defect within 250 ms of detecting the onset of at least five consecutive samples (which may or may not be consecutive frames) of matched STS Signal Labels, as specified in Table 6-2.

For testing conformance to this requirement, it is sufficient to transmit a continuous string of matched STS Signal Labels for at least 250 ms.

- O6-127** [477] STS PTE should terminate a PLM-P defect immediately upon receipt of five contiguous frames with matched STS Signal Labels, as specified in Table 6-2.
- R6-128** [478v2] STS PTE shall terminate a PLM-P defect upon detecting an UNEQ-P defect.
- R6-129** [479] An NE shall declare a PLM-P failure if a PLM-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a PLM-P failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
- R6-130** [480v2] An NE shall clear the PLM-P failure if the PLM-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the PLM-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.1.8.B *STS Path Unequipped*

- R6-131** [481] STS PTE shall detect an STS Path Unequipped (UNEQ-P) defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of unequipped STS Signal Labels (C2 byte), as specified in Table 6-2.
- O6-132** [482] STS PTE should detect an UNEQ-P defect immediately upon receipt of five contiguous frames with unequipped STS Signal Labels, as specified in Table 6-2.
- R6-133** [485] STS PTE shall terminate an UNEQ-P defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of STS Signal Labels that are not unequipped, as specified in Table 6-2.
- O6-134** [486] STS PTE should terminate an UNEQ-P defect immediately upon receipt of five contiguous frames with STS Signal Labels that are not unequipped, as specified in Table 6-2.
- R6-135** [488] An NE shall declare an UNEQ-P failure if an UNEQ-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an UNEQ-P failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.

- R6-136** [489v2] An NE shall clear the UNEQ-P failure if the UNEQ-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the UNEQ-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

When a service is first activated, it is possible that STS PTE would detect an UNEQ-P defect until the service is properly and completely provisioned and activated. As an option, to avoid declaring and reporting UNEQ-P failures when a service is first activated, a SONET NE may be designed such that it does not detect an UNEQ-P defect until the service is activated.

Table 6-2. STS Signal Label Mismatch Defect Conditions

Provisioned STS PTE Functionality	Received Payload Label (C2 Byte, in hex format)	Defect
None ^a	Anything	None
Any Equipped functionality ^b	Unequipped (00)	UNEQ-P
Any Equipped functionality ^b	Equipped – Non Specific (01)	None (Matched)
Equipped – Non Specific	A value corresponding to any Payload Specific functionality (02 to E0, or FD to FE)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to the same Payload Specific functionality as the provisioned functionality (02 to E0, or FD to FE)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to a different Payload Specific functionality than the provisioned functionality (02 to E0, or FD to FE)	PLM-P
Equipped – Non Specific, or VT-Structured STS-1	PDI, 1 to 27 VTx defects (E1 to FB)	None (Matched)
Any Payload Specific functionality except VT-Structured STS-1 ^b	PDI, 1 to 27 VTx defects (E1 to FB)	PLM-P
Any Equipped functionality ^b	PDI, 28 VT1.5 defects or 1 non-VT payload defect (FC)	None (Matched)
Any Equipped functionality ^b	Reserved ^c (FF)	No change ^d

Notes:

- This entry corresponds to the case when an NE is not provisioned to expect a signal. See Section 3.3.2.3.
- See Table 3-2 for the currently assigned STS signal label values.
- The all-ones signal label is reserved in ITU-T G.707 for an AIS function (and is not currently defined for any use in this document). It is not expected to be assigned to a specific payload.
- That is, the detection of 'FF' in the STS signal label shall not cause a new UNEQ-P or PLM-P defect to be detected for that path, or cause an existing UNEQ-P or PLM-P defect to be terminated.

6.2.1.1.8.C VT Payload Label Mismatch

R6-137 [491] VT PTE shall detect a VT Payload Label Mismatch (PLM-V) defect within 250 ms of the onset of at least five consecutive samples

(which may or may not be consecutive superframes) of mismatched VT Signal Labels, as specified in Table 6-3.

For testing conformance to this requirement, it is sufficient to transmit a continuous string of identical, mismatched VT Signal Labels in bits 5 through 7 of the V5 byte for at least 250 ms.

O6-138 [492] VT PTE should detect a PLM-V defect immediately upon receipt of five contiguous superframes with mismatched VT Signal Labels, as specified in Table 6-3.

R6-139 [495] VT PTE shall terminate a PLM-V defect within 250 ms of detecting the onset of at least five consecutive samples (which may or may not be consecutive superframes) of matched VT Signal Labels, as specified in Table 6-3.

For testing conformance to this requirement, it suffices to transmit a continuous string of matched VT Signal Labels for at least 250 ms.

O6-140 [496] VT PTE should terminate a PLM-V defect immediately upon receipt of five contiguous superframes with matched VT Signal Labels, as specified in Table 6-3.

R6-141 [497v2] VT PTE shall terminate the PLM-V defect immediately upon detecting an UNEQ-V defect on the incoming signal.

R6-142 [498] An NE shall declare a PLM-V failure when a PLM-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a PLM-V failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.

R6-143 [499] An NE shall clear a PLM-V failure when the PLM-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the PLM-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

6.2.1.1.8.D *VT Path Unequipped*

R6-144 [501] VT PTE shall detect a VT Path Unequipped (UNEQ-V) defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive superframes) of unequipped VT Signal Labels, as specified in Table 6-3.

- O6-145** [502] VT PTE should detect an UNEQ-V defect immediately upon receipt of five contiguous superframes with unequipped VT Signal Labels, as specified in Table 6-3.
- R6-146** [505] VT PTE shall terminate an UNEQ-V defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive superframes) of VT Signal Labels that are not unequipped, as specified in Table 6-3.
- O6-147** [506] VT PTE should terminate an UNEQ-V defect immediately upon receipt of five contiguous superframes with VT Signal Labels that are not unequipped, as specified in Table 6-3.
- R6-148** [508] An NE shall declare an UNEQ-V failure when an UNEQ-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an UNEQ-V failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
- R6-149** [509] An NE shall clear an UNEQ-V failure when the UNEQ-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the UNEQ-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

When a service is first activated, it is possible that VT PTE would detect an UNEQ-V defect until the service is properly and completely provisioned and activated. As an option, to avoid declaring and reporting UNEQ-V failures when a service is first activated, a SONET NE may be designed such that it does not detect an UNEQ-V defect until the service is activated.

Table 6-3. VT Signal Label Mismatch Defect Conditions

Provisioned VT PTE Functionality	Received Payload Label (V5 Bits 5-7)	Defect
None (zero) ^a	Anything	None
Any Equipped functionality ^b	Unequipped (000)	UNEQ-V
Any Equipped functionality ^b	Equipped – Non Specific (001)	None (Matched)
Equipped – Non Specific	A value corresponding to any Payload Specific functionality (010 to 110)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to the same Payload Specific functionality as the provisioned functionality (010 to 110)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to a different Payload Specific functionality than the provisioned functionality (010 to 110)	PLM-V
Any Equipped functionality ^b	Reserved ^c (111)	No change ^d

Notes:

- This entry corresponds to the case when an NE is not provisioned to expect a signal. See Section 3.3.3.
- See Table 3-4 for the currently assigned VT signal label values.
- The all-ones signal label is reserved in ITU-T G.707 (but not currently defined in this document) for an AIS function. It is not expected to be assigned to a specific payload.
- That is, the detection of '111' in the VT signal label shall not cause a new UNEQ-V or PLM-V defect to be detected for that path, or cause an existing UNEQ-V or PLM-V defect to be terminated.

6.2.1.1.9 Path Trace Identifier Mismatch

This section contains information and criteria related to the use of the STS path trace string (i.e., the J1 byte) for STS Path Trace Identifier Mismatch (TIM-P) defect/failure detection purposes. Similar to an UNEQ defect, a TIM defect is considered a connectivity defect because it is expected to be caused by provisioning problems (e.g., incorrect cross-connections) within the network. In the case of a TIM defect, the incorrect provisioning causes PTE at one NE to be connected through the network to the wrong far-end PTE. At this time, TIM defects and failures have been defined in SONET only for STS paths; however, they could be defined for VT paths in the future. In addition, in order to coexist with equipment that does not support the STS path trace provisioning capabilities

necessary for reliable TIM-P detection purposes (e.g., equipment that does not meet R6-151 [997v2] and R6-152 [1069]), an NE that supports the TIM-P detection feature must be capable of having that feature activated and deactivated on a per-STS path basis.

Note that prior to GR-253-CORE, Issue 2, Revision 2, a received STS path trace string was expected to be monitored and used for diagnostics purposes only. When used for diagnostics purposes the received path trace string may be compared to an "expected" path trace string (as is done for TIM detection purposes), and persistent mismatches may be reported to an OS. However, the detection of a mismatch on a path trace string that is being monitored for diagnostics purposes does not cause the generation of AIS downstream or RDI upstream, and does not affect the accumulation of any PM parameters. In general (and assuming ERDI is supported), all of those actions are currently required to be taken when a TIM defect is detected (see Sections 6.2.1.2.3, 6.2.1.2.4, 6.2.1.3.2, and 6.2.2.5).⁵

R6-150 [725v2] A SONET NE that contains STS PTE shall allow the user to provision, on a per-path basis, the contents of the STS Path Trace carried in the J1 byte of the STS path overhead originated by the PTE. The transmitted STS Path Trace string shall be 64 bytes in length.

In most situations it is assumed that the user will provision STS PTE in a SONET NE to transmit path trace strings that consist of ASCII printable characters [i.e., '20' through '7E' (hex)], and the following criteria are generally based on that assumption. However, in the ITU-T SDH standards the J1 byte may carry a 16-byte identifier⁶ rather than a 64-byte string. The criteria in this section are not intended to preclude a SONET NE from supporting the capability to transmit or detect path trace strings that meet the SDH 16-byte format (e.g., for use at international boundaries).

R6-151 [997v2] A SONET NE that contains STS PTE shall support a feature that allows the contents of the STS Path Traces to be provisioned as ASCII characters. In addition, the following apply:

- ... • The feature shall allow the user to enter a string of up to 62 characters
- ... • The feature shall place no restriction on the content of the string, except that the characters shall be ASCII printable characters

5. Note that at the time that this section was added to GR-253-CORE, the insertion of AIS downstream upon detection of a TIM-P defect was under study in ANSI T1X1.5. There it had been suggested that instead of (or possibly in addition to) requiring that the TIM-P defect detection feature be a provisionable option (see R6-153 [1070]), there should be a requirement to state that the insertion of AIS in response to a TIM-P defect must be provisionable. This topic is discussed further in GR-253-ILR Issue ID 253-139.

6. The format for the 16-byte path trace is described in ITU-T G.831, *Management capabilities of transport networks based on the Synchronous Digital Hierarchy (SDH)*, as a frame start marker, the result of a CRC-7 calculation over the previous frame, and 15 ITU-T Recommendation T.50 characters beginning with either the country code as defined in ITU-T E.164 or the alphabetic character country code as defined in ISO 3166. The T.50 and ASCII character sets are essentially the same, so 15 bytes of the 16-byte path trace could be considered "ASCII-based". However, the byte containing the frame start marker and the CRC-7 result cannot be considered "ASCII-based".

- ...
- The NE shall automatically pad the string entered by the user to 62 characters using ASCII NULL characters, and then add <CR> and <LF> characters (i.e., '0D' and '0A') for a total of 64 characters.
- ...
- Each 8-bit ASCII character shall be loaded into one J1 byte.

For example, a user-provisioned ASCII CLLI code, padded with NULL characters and terminated with <CR> and <LF> characters (for a total of 64 bytes) would be a suitable path trace message.

As discussed above for the transmitted STS Path Trace, it is assumed that the path trace strings detected by SONET NEs for the purpose of TIM-P detection will generally consist of ASCII printable, NULL, <CR>, and <LF> characters. The following requirement is based on that assumption.

R6-152 [1069] A SONET NE shall support a feature to allow the user to provision the "expected" ASCII-based path trace for each STS path that it terminates and for which TIM-P detection has been activated (see below). In addition, the following apply:

- ...
- The feature shall allow the user to enter a string of up to 62 characters
- ...
- The feature shall place no restriction on the contents of the string, except that the characters shall be ASCII printable characters.

The capability to provision an expected STS path trace may also need to be provided by an NE that supports the STS path trace diagnostics described in Section 6.2.3.2.3.A (see **CR6-398 [999v2]**). Therefore it is likely that an NE will need to provide the capability to activate and deactivate TIM-P detection (see below) separately from the provisioning of the expected trace (e.g., when an expected path trace string is entered, the user may also need to provision whether it is to be used for TIM-P detection or diagnostics purposes).

R6-153 [1070] A SONET NE that contains STS PTE shall support a feature that, when activated monitors the received STS path trace string for TIM-P detection purposes. That feature shall be provisionable on a per-STs path basis, and shall have a default of "not active".

In general, the possible impacts (both positive and negative) should be carefully considered before the detection of TIM-P defects is activated for a particular path. For example, the user should be aware that due to interactions between various criteria (including the current PM parameter definitions in Section 6.2.2), the activation of TIM-P defect detection for a path could result in discrepancies between different sets of PM data accumulated for that path.⁷

Since TIM defects are expected to occur primarily as a result of provisioning errors (e.g., incorrect cross-connections in the network), they are generally expected to be of relatively long duration, and it is not considered essential that they be detected and terminated as quickly as other types of defects. Thus, T1.231-1997 allows up to 30 seconds for the

detection or termination of a TIM-P defect. This is reflected in **R6-154 [1071]** and **R6-159 [1075]**; however as indicated in **O6-158 [1074]** and **O6-160 [1076]**, significantly faster defect detection and termination times are desirable.

R6-154 [1071] STS PTE shall detect a TIM-P defect within 30 seconds (or less) when none of the sampled 64-byte STS path trace strings match the provisioned expected value.

Note that if the STS PTE has detected an AIS-P or LOP-P defect on its incoming signal, the J1 byte cannot be accessed to monitor for TIM-P defects, and therefore a TIM-P defect cannot be either detected or terminated. Otherwise, during a 30-second interval the STS PTE can expect to receive 3750 copies of the incoming path trace string. To reduce the processing burden on the NE, T1.231-1997 and **R6-154 [1071]** indicate (via the use of the word "sampled") that not all of these copies need to be compared to the provisioned expected value. While no specific value is provided for the number of consecutive non-matching samples that should trigger the detection of a TIM defect, that number needs to be large enough to avoid "false" TIM defect detection (i.e., the detection of TIM defects that are caused by random bit errors rather than an actual end-to-end connectivity problem).

R6-155 [1072] A SONET NE's TIM-P defect detection algorithm shall be such that, given an incoming signal with a BER of 10^{-3} or less, the probability that the STS PTE will detect a false TIM-P defect during the TIM-P defect detection time is less than 1×10^{-n} .

The current value for "n" in **R6-155 [1072]** is 6; however, this value is under study and may need to be changed in the future. Also, the "TIM-P defect detection time" is the time that it would take the STS PTE to detect a TIM-P defect if the incoming path trace were to change to a new (non-matching) value. For example, this could be the path trace string sampling rate times the number of non-matching samples necessary to cause the detection of a TIM-P defect.

In addition to random bit errors, an NE's TIM defect detection algorithm needs to be able to tolerate changes in the phase of the incoming path trace string. Such changes could be caused, for example, by protection switches at upstream NEs that result in one or more J1 bytes being dropped or repeated (e.g., due to transmission delay differences between the long and short paths around a UPSR).

7. As noted previously, if ERDI-P is supported for a particular STS path, then TIM-P defects/failures are currently required to affect the accumulation of certain PM parameters. For a particular STS path this effect is both direct (i.e., based on the presence or absence of the TIM-P defect or failure) at the STS PTE that is accumulating the near-end STS path PM data, and indirect (e.g., via the detection of ERDI-P Connectivity defects) at any upstream NEs that are accumulating the corresponding far-end STS path or intermediate-path PM data. On the other hand, any intermediate NEs that are accumulating near-end intermediate-path PM involving that path are currently not expected to be able to monitor for TIM-P defects/failures. This could lead to significant discrepancies between the various sets of PM data accumulated for the path, and therefore it has been suggested that it may be desirable to revisit some of the relevant criteria (e.g., the criteria that define the TIM-P defect, the ERDI-P trigger criteria, the PM parameter definitions). This issue is for further study.

R6-156 [1073] A change in the phase of an incoming STS path trace string shall cause the STS PTE to consider, at most, one sample to be mismatched for the purpose of detecting and terminating TIM-P defects.

O6-157 [1001v2] If a SONET NE is comparing an incoming path trace string to an expected path trace (for either TIM-P defect detection/termination or diagnostics purposes) and the expected path trace consists of ASCII characters, then the comparison should ignore any trailing ASCII NULL, <CR>, and <LF> characters contained in the incoming path trace.

An NE that meets **O6-157** [1001v2] could accept (and consider matched) an incoming path trace generated by STS PTE that meets **R6-151** [997v2] (i.e., that pads with NULLs to 62 characters and then adds <CR> and <LF>), or by older STS PTE that might not meet **R6-151** [997v2] (e.g., that pads with NULLs to 64 characters, or that adds the <CR> and <LF> characters before padding with NULLs).

O6-158 [1074] A SONET NE's STS path trace sampling rate and TIM-P defect detection algorithm should be such that a TIM-P defect is detected within 2 seconds after the NE's STS PTE stops receiving the provisioned expected path trace string.

R6-159 [1075] STS PTE shall terminate a TIM-P defect within 30 seconds (or less) when four-fifths (or more) of the sampled STS path trace strings match the provisioned expected value.

Note that in order to meet **R6-159** [1075] it would be sufficient for STS PTE to sample the incoming STS path trace string as infrequently as once every six seconds (i.e., 5 samples in any 30-second window). However, such a sampling rate would not be sufficient for detecting TIM-P defects as it would result in an unacceptably high probability of false TIM-P defect detection.

O6-160 [1076] A SONET NE's STS path trace sampling rate and TIM-P defect termination algorithm should be such that a TIM-P defect is terminated within 2 seconds after the STS PTE begins receiving the provisioned expected path trace string.

R6-161 [1077] An NE shall declare a TIM-P failure if a TIM-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a TIM-P failure indication and send an alarm message to an OS unless the condition in **R6-285** [626v2] applies.

Unlike most other cases involving failures based on directly detected defects, it is very likely that the conditions for declaring both a TIM-P failure and either an UNEQ-P or PLM-P failure will be met in response to a single root-cause incoming signal problem. That is, an incorrect cross-connection in the network is likely to result in the STS PTE

receiving either an unequipped STS or an STS containing the wrong payload mapping, and also an STS with a path trace string that does not match the provisioned expected value. The TIM-P defect has not been defined such that its detection automatically causes the termination of either an UNEQ-P or PLM-P defect (or vice versa), and therefore the NE's conformance to **R6-285 [626v2]** is particularly important to avoid the generation of extraneous messages to the OS. However, if the NE does not meet **O6-158 [1074]** then it is also very likely that the TIM-P defect will not yet have been detected at the time when the condition for declaring the UNEQ-P or PLM-P failure is met. For an UNEQ-P failure this is not an issue since that failure is considered higher priority than a TIM-P failure (see Table 6-6). That is, when the TIM-P failure is subsequently declared it would simply not be autonomously reported to the OS. On the other hand, it is an issue for a PLM-P failure since that failure is considered lower priority than a TIM-P failure. In that case the NE may or may not clear the PLM-P failure when the TIM-P failure is subsequently declared.

- R6-162 [1078]** An NE shall clear the TIM-P failure if the TIM-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the TIM-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

Note that since the maximum TIM-P defect termination time of 30 seconds is much longer than the failure-declaration soaking time of approximately 2.5 seconds, there do not appear to be any situations in which an NE would be required to detect and terminate a TIM-P defect without also declaring a TIM-P failure. Similarly, since the maximum TIM-P defect detection time of 30 seconds is much longer than the failure-clearing soaking time of approximately 10 seconds, there do not appear to be any situations in which an NE would be required to terminate and (re)detect a TIM-P defect without also clearing a TIM-P failure. However, the 2.5 and 10 seconds soaking intervals are included in **R6-161 [1077]** and **R6-162 [1078]** (and in the corresponding T1.231-1997 specifications) for consistency with the requirements for other types of failures. In addition, those soaking intervals could be of importance for NEs that detect and terminate TIM-P defects substantially faster than the maximum allowed detection and termination times of 30 seconds.

6.2.1.2 Alarm Indication Signal

An Alarm Indication Signal (AIS) is a maintenance signal used in the digital network to alert downstream equipment that a defect or equipment failure has been detected. The SONET signal format provides AISs for the Line (AIS-L), STS Path (AIS-P), and VT Path (AIS-V) layers. In addition, SONET NEs may also deal with DS_n (e.g., DS0, DS1, DS3) AIS and trunk conditioning.

In most cases, equipment generates the AIS defined for the next higher layer when it detects a defect on the signal that it is terminating (e.g., STE generates AIS-L when it detects an LOS or LOF defect on its incoming signal), as depicted in Figure 6-2. Figure 6-2 also illustrates the upstream maintenance signals discussed in Section 6.2.1.3, and shows that

AIS is sent downstream in both the physical sense (i.e., as an externally discernible signal) and the logical sense (i.e., as an internally generated signal communicated between logical SONET layers). Thus, a defect detected at a given layer (e.g., Section layer) functionally causes an internal AIS to be generated to all applicable higher layers (e.g., AIS-L).

In addition to being generated by equipment that detects a defect on the signal that it is terminating, in some cases AIS is also generated by equipment that originates a signal. This occurs when equipment that is originating a signal detects that the circuitry supporting provisioned higher layer origination functions has failed or been removed without having been "unprovisioned", and it continues until standby circuitry (if available) is switched in or the failure is cleared. For example, LTE generates AIS-P (for the affected STS paths) when it detects that STS PTE supporting provisioned STS path origination functions has failed.

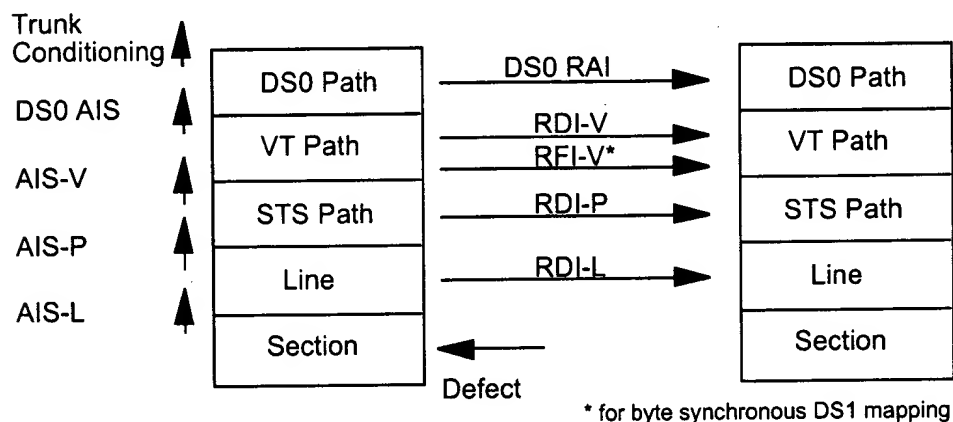


Figure 6-2. Maintenance Signals for SONET Layers

6.2.1.2.1 Line AIS (AIS-L)

STE (and physical layer regenerators, see TR-NWT-000917) send AIS-L to alert the downstream LTE that a defect has been detected on the incoming SONET section, or that LTE supporting provisioned line origination functions has failed. If line-level APS is provided as a feature, the downstream LTE can initiate protection switching upon detection of the AIS-L defect, as specified in Section 5.3.3.1 for linear APS.

- R6-163** [512v2] STE shall generate AIS-L downstream within 125 μ s of detecting an LOS or LOF defect on the incoming signal or the failure of LTE supporting provisioned line origination functions. The AIS-L shall be generated as an OC-N or STS-N electrical signal that contains valid

Section overhead and a scrambled all-ones pattern for the remainder of the signal.

Note that the AIS-L generated as described above automatically provides convenient generation of AIS for higher layers (e.g., STS and VT Path AIS). Figures 6-3 and 6-14 (in Section 6.2.1.7) illustrate when STE generates AIS-L in response to a defect detected on an incoming signal or an equipment failure.

- R6-164** [513v2] STE shall deactivate AIS-L within 125 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 125 μ s of clearing the failure or determining that standby equipment has been switched in.
- R6-165** [514] LTE shall detect an AIS-L defect on the incoming signal when bits 6, 7, and 8 of the K2 byte contain the '111' pattern in five consecutive frames.
- R6-166** [515] LTE shall terminate the AIS-L defect on the incoming signal when bits 6, 7, and 8 of the K2 byte have any pattern other than '111' in five consecutive frames.

For testing conformance to **R6-166 [515]**, it is sufficient to use a consistent non-'111' pattern.

- R6-167** [516] An NE shall declare an AIS-L failure if an AIS-L defect persists for 2.5 (± 0.5) seconds. Upon declaring an AIS-L failure, the NE shall set an AIS-L failure indication for the line and (if AIS-L is a Reported failure for the line) send a message to an OS unless the condition in **R6-285 [626v2]** applies.
- R6-168** [517] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare AIS-L failures. Upon declaring an AIS-L failure, the NE shall perform the actions listed in **R6-167 [516]**.
- R6-169** [518] An NE shall clear the AIS-L failure when the AIS-L defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-L failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.2.2 STS Path AIS (AIS-P)

LTE sends AIS-P to alert the downstream STS PTE that it has detected a defect on its incoming line signal, or that STS PTE supporting provisioned path origination functions has failed. Figures 6-4 and 6-14 (in Section 6.2.1.7) illustrate when AIS-P is generated in response to a defect detected on an incoming signal or an equipment failure.

- R6-170** [519v3] LTE shall generate AIS-P downstream for the affected STS paths within 125 μ s of detecting an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) or (if the STS pointer is processed) an LOP-P defect on the incoming signal, or the failure of STS PTE supporting provisioned path origination functions. The AIS-P shall be generated as all-ones in the H1, H2 and H3 bytes, and the entire STS SPE.

Note that for an STS-Mc path, all-ones is generated in all M of the H1, H2 and H3 bytes. Also note that in the functional model used in this document, incoming defects that are detected at STE (i.e., LOS, LOF) cause AIS-L to be generated downstream. This AIS-L is then detected by the LTE and causes it generate AIS-P.

- R6-171** [521v2] LTE shall deactivate the AIS-P within 125 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 125 μ s of clearing the failure or determining that standby equipment has been switched in. LTE that performs STS pointer processing shall deactivate AIS-P by constructing a correct STS pointer with a set NDF, followed by normal pointer operations, as well as ceasing to insert the all-ones pattern in the STS SPE. LTE that does not perform STS pointer processing shall deactivate AIS-P by ceasing the insertion of all-ones in the H1, H2 and H3 bytes, and the STS SPE.

- R6-172** [523] STS PTE shall detect an AIS-P defect when the H1 and H2 bytes for an STS path contain an all-ones pattern in three consecutive frames. For an STS-Nc path, only the H1 and H2 bytes of the first STS-1 need to be observed.

Note that in the functional model used in this document, incoming AIS-P defects are detected by STS PTE, not by LTE. If LTE that processes STS pointers receives all-ones H1 and H2 bytes, it is required to relay those all-ones values downstream (see Section 3.5.1.5). No "all-ones pointer relay" defect has been defined (and therefore an all-ones pointer relay failure is not declared if the LTE continues to relay the all-ones pointer); however, the act of performing all-ones pointer relay does have several effects on the operation of LTE. For example, it affects the LTE's detection of LOP-P defects (see Section 6.2.1.1.3). In addition it should be noted that since LTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, the incoming STS SPE associated with the all-ones pointers cannot be expected to be passed through

unaltered. As a minimum, it is likely that the SPE sent downstream will be affected by buffer overflows or underflows that occur at the LTE. Alternatively, some LTE may discard the incoming SPE and generate a new all-ones SPE (synchronized to the local clock) for transmission downstream.

R6-173 [524] STS PTE shall terminate an AIS-P defect when the H1 and H2 bytes for the STS path contain a valid STS Pointer with a set NDF, or when they contain valid, identical STS Pointers with normal NDFs for three consecutive frames. For an STS-Nc path, the concatenation indicators must also be valid.

O6-174 [1079] STS PTE should terminate an AIS-P defect when it detects an LOP-P defect.

Note that prior to Issue 2, Revision 2 of this document, STS PTE was required to terminate an AIS-P defect only in the cases shown in **R6-173** [524], and to *not* detect an LOP-P defect if an AIS-P defect had already been detected. Therefore, NEs designed to earlier issues of this document are not expected to conform to **O6-174** [1079]. Also note that the current criteria are consistent with the specifications that appear in ITU-T Recommendation G.783, while the previous criteria were consistent with those that appeared in ANSI T1.231-1997. At the time that Issue 2, Revision 2 was released, the specifications in T1.231 were under review and were expected to be modified to also be consistent with the G.783 specifications.

R6-175 [525] An NE shall declare an AIS-P failure if an AIS-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an AIS-P failure, the NE shall set an AIS-P failure indication for that path and (if AIS-P is a Reported failure for the path) send a message to an OS unless the condition in **R6-285** [626v2] applies.

R6-176 [526] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56** [419] to declare AIS-P failures. Upon declaring an AIS-P failure, the NE shall perform the actions listed in **R6-175** [525].

R6-177 [527] An NE shall clear an AIS-P failure when the AIS-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.2.3 VT Path AIS (AIS-V)

STS PTE sends AIS-V to alert the downstream VT PTE that it has detected a defect on its incoming path signal, or that VT PTE supporting provisioned path origination functions has failed. In addition, VT PTE that byte-synchronously maps an incoming DS1 into a VT1.5 sends AIS-V when certain DS1 defects are detected. Figures 6-6, 6-10, and 6-14 (in Section 6.2.1.7) illustrate when AIS-V is generated in response to a defect detected on an incoming signal or an equipment failure.

- R6-178** [528v4] STS PTE shall generate AIS-V downstream for the affected VT paths within 500 μ s of detecting an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-P defect, an UNEQ-P defect, a TIM-P defect (if activated, also see GR-253-ILR Issue ID 253-139), a PLM-P defect, or (if the VT pointer is processed) an LOP-V defect on the incoming signal, or the failure of VT PTE supporting provisioned path origination functions. The AIS-V shall be generated as an all-ones code in the entire VT, including the V1 through V4 bytes.

Note that in the functional model used in this document, incoming defects that are detected at STE or LTE (i.e., LOS, LOF, AIS-L) cause AIS-P to be generated downstream. This AIS-P is then detected by the STS PTE and causes it generate AIS-V.

- R6-179** [529] VT PTE shall generate AIS-V within 500 μ s of detecting a DS1 LOS, OOF, or AIS defect on an incoming DS1 that it byte-synchronously maps into a single VT1.5.
- R6-180** [531v2] STS PTE shall deactivate AIS-V within 500 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 500 μ s of clearing the failure or determining that standby equipment has been switched in. STS PTE that performs VT pointer processing shall deactivate AIS-V by constructing a correct VT pointer with valid VT size and a set NDF, followed by normal pointer operations, as well as ceasing to insert the all-ones pattern in the rest of the VT. STS PTE that does not performing VT pointer processing shall deactivate AIS-V by ceasing the insertion of the all-ones pattern in the entire VT.
- R6-181** [533] VT PTE shall deactivate AIS-V (as described in **R6-180** [531v2]) within 500 μ s of terminating the defect on the incoming DS1 that caused it to be generated.
- R6-182** [534] VT PTE shall detect an AIS-V defect for the VT path upon receiving an all-ones pattern in the V1 and V2 bytes in three consecutive superframes.

Note that in the functional model used in this document, incoming AIS-V defects are detected by VT PTE, not by STS PTE. If STS PTE that processes VT pointers receives all-ones V1 and V2 bytes, it is required to relay those all-ones values downstream (see Section 3.5.2.5). No "all-ones pointer relay" defect has been defined (and therefore an all-ones pointer relay failure is not declared if the STS PTE continues to relay the all-ones pointer); however, the act of performing all-ones pointer relay does have several effects on the operation of STS PTE (see Section 6.2.1.1.3). In addition it should be noted that since STS PTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, the incoming VT SPE associated with the all-ones pointers cannot be expected to be passed through unaltered. As a minimum, it is likely that the SPE sent downstream will be affected by buffer overflows or underflows that occur at the STS PTE. Alternatively, some STS PTE may discard the incoming SPE and generate a new all-ones SPE (synchronized to the local clock) for transmission downstream.

R6-183 [535] VT PTE shall terminate an AIS-V defect upon receiving a valid VT Pointer with valid VT size and a set NDF, or upon receiving three consecutive superframes containing valid, identical VT Pointers with a valid VT size and normal NDFs.

O6-184 [1080] VT PTE should terminate an AIS-V defect when it detects an LOP-V defect.

Note that prior to Issue 2, Revision 2 of this document, VT PTE was required to terminate an AIS-V defect only in the cases shown in **R6-183** [535], and to *not* detect an LOP-V defect if an AIS-V defect had already been detected. Therefore, NEs designed to earlier issues of this document are not expected to conform to **O6-184** [1080]. Also note that the current criteria are consistent with the specifications that appear in ITU-T Recommendation G.783, while the previous criteria were consistent with those that appeared in ANSI T1.231-1997. At the time that Issue 2, Revision 2 was released, the specifications in T1.231 were under review and were expected to be modified to also be consistent with the G.783 specifications.

R6-185 [536] An NE shall declare an AIS-V failure if an AIS-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the AIS-V failure, the NE shall set an AIS-V failure indication for the path and (if AIS-V is a Reported failure for the path) send a message to an OS unless the condition in **R6-285** [626v2] applies.

R6-186 [537] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56** [419] to declare AIS-V failures. Upon declaring an AIS-V failure, the NE shall perform the actions listed in **R6-185** [536].

- R6-187** [538] An NE shall clear an AIS-V failure when the AIS-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-V failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.2.4 DS_n AIS

In many applications, SONET is the transport medium for lower-speed digital signals such as DS1s and DS3s, and in those applications the SONET NEs may need to generate or detect DS_n AIS. Figures 6-5 through 6-12 (in Section 6.2.1.7) illustrate the use of DS_n AIS for STS and VT PTE terminating SPEs of various compositions. Criteria on the construction of DS1, DS1C, DS2, and DS3 AIS are contained in GR-499-CORE. DS0 AIS (also sometimes referred to as "UNICODE") is defined in SONET as a specific pattern in the ABCD signaling bits (i.e., A = 0, B = 0, C = 1, D = 0). Note that DS0 AIS may not be applicable outside of SONET networks, so an interface to a non-SONET NE may require the application of service-specific trunk conditioning codes (i.e., one or more provisioned ABCD signaling codes and an 8-bit insertion word). Also note that DS0 AIS is defined to include an all-ones pattern in the DS0 payload (in addition to the '0010' code in the signaling bits) in GR-303-CORE, *Integrated Digital Loop Carrier System Generic Requirements, Objectives, and Interface*. For defects detected on an incoming SONET signal, the generation of a lower-layer AIS (e.g., AIS-P, AIS-V) could be sufficient to ensure all-ones DS0 payloads. However, for defects detected on an incoming DS1, it may be necessary for the NE to actively insert all-ones in the affected DS0s.

Criteria on generating AIS at other transport signal interfaces is for further study.

- R6-188** [539v2] STS or VT PTE shall generate DS1, DS1C, DS2, or DS3 AIS downstream within 125 μ s of the detection of certain defects, as shown in Figures 6-5 through 6-10.
- R6-189** [951] VT PTE with DS0 rearrangement capabilities shall generate DS0 AIS downstream within 3 ms of the detection of certain defects (unless it is provisioned to apply a service-specific trunk conditioning code, see Section 6.2.1.6), as shown in Figures 6-11 and 6-12.

Since DS0 AIS is carried in the ABCD signaling bits, which have a cycle time of 3 ms, VT PTE could meet the preceding requirement by (for example) sending the DS0 AIS code in the next complete set of ABCD signaling bits after the defect is detected.

- R6-190** [540] STS or VT PTE shall remove a downstream DS1, DS1C, DS2, or DS3 AIS within 125 μ s of terminating the defect that caused it to be sent.
- R6-191** [541v2] If a defect that causes VT PTE to insert DS0 AIS is terminated before a failure is declared, then the VT PTE shall remove the DS0 AIS

within 3 ms of terminating that defect. If a failure was declared, then the VT PTE shall remove the DS0 AIS within 3 ms of clearing the failure.

In addition to generating and removing DS_n AIS, in some applications a SONET NE may also need to be able to (non-intrusively) monitor for DS_n AIS even if it does not terminate the DS_n path. For example, STS PTE that supports the generation of PDI-P signals and uses the asynchronous DS3 mapping would need to be capable of detecting DS3 AIS on the incoming DS3 signal in order to insert the appropriate code in the STS signal label. Similarly, an NE that performs intermediate DS_n Path PM would need to monitor for DS_n AIS to correctly accumulate certain parameters.

R6-192 [542] A SONET NE shall monitor for DS_n AIS and declare DS_n AIS failures on DS_n paths that it terminates.

Requirements on DS_n AIS detection and failure declaration for DS1, DS1C, DS2, and DS3 signals are contained in GR-499-CORE. For DS0 AIS, see **R6-196** [547] through **R6-199** [955].

CR6-193 [544v2] A SONET NE with DS_n interfaces may be required to monitor for DS_n AIS (as if it were terminating the DS_n path) on incoming DS_n signals where the DS_n path is not terminated.

CR6-194 [952] A SONET NE with DS_n interfaces may be required to monitor for DS_n AIS (as if it were terminating the DS_n path) on DS_n signals that are asynchronously demultiplexed from STS or VT SPEs (i.e., on outgoing DS_n signals).

See Section 6.2.1.8.7 for criteria on the declaration of DS_n AIS failures when the DS_n path is not terminated.

O6-195 [953] If the DS_n path is not terminated, then the capability to monitor for DS_n AIS should be provided independently of any options to provide non-clear-channel transport of DS_n signals using the asynchronous DS_n mappings (see Section 6.2.1.1.2).

An NE that meets **O6-195** [953] would be capable of simultaneously providing (for example) both clear-channel DS3 transport and PDI-P generation for the STS-1 SPE containing that DS3. Note that in order to detect DS_n AIS, an NE may need to monitor the framing bits in the DS_n (e.g., to detect DS1 AIS an NE must determine that valid DS1 framing is not present, and to detect DS3 AIS it must determine that valid DS3 framing is present). However, an NE should be able to perform that monitoring without also inserting DS_n AIS downstream when a DS_n OOF is detected.

R6-196 [547] A SONET NE that detects DS0 AIS shall detect a DS0 AIS defect if it receives two consecutive sets of ABCD signaling bits set to the code '0010' (i.e., the '0010' code persisting for 6 ms).

- R6-197** [548] A SONET NE shall terminate a DS0 AIS defect if it receives four consecutive sets of ABCD signaling bits set to any code other than the '0010' code (i.e., any code other than '0010' persisting for 12 ms).
- R6-198** [954] A SONET NE that is provisioned to insert a service-specific trunk conditioning code on a particular DS0 path shall declare a DS0 AIS failure when a DS0 AIS defect persists for 3.25 (± 0.25) seconds.
- R6-199** [955] A SONET NE that is provisioned to insert a service-specific trunk conditioning code on a particular DS0 path shall clear a DS0 AIS failure when the DS0 AIS defect is terminated.

See Section 6.2.1.6 for criteria related to the actions that an NE provisioned to insert a service-specific trunk conditioning code takes when it declares or clears a DS0 AIS failure. Criteria on the declaration and clearing of DS0 AIS failures by other NEs that detect DS0 AIS are provided in the GRs, TRs, or TAs for those NEs (e.g., GR-303-CORE for IDLC NEs).

6.2.1.3 Remote Defect Indication (RDI) and Remote Failure Indication (RFI)

Remote Alarm Indication (RAI) signals have historically been used in the digital network to alert upstream terminals of a downstream failure so that they can initiate trunk conditioning on the failed circuit. The layered maintenance strategy incorporated in the SONET signal format (as shown in Figure 6-2) expands the applications of these signals in the network. In SONET, RDI signals occur at the Line, STS Path, and VT Path layers, and a persistent RDI defect results in the derivation of an RFI failure. In addition, for byte-synchronously mapped DS1s, a VT Path layer RFI signal is also defined, both to provide backward compatibility with equipment using the traditional RAI signal timing, and for translation to and from DS1 RAI signals.

6.2.1.3.1 Line Remote Defect Indication (RDI-L) and Remote Failure Indication (RFI-L)

The RDI-L signal (formerly called Line FERF) indicates to LTE that its peer LTE has detected an AIS-L (or lower layer) defect on the signal that it (the first LTE) originated. An incoming RDI-L defect is used to derive an RFI-L failure.

- R6-200** [549v2] LTE shall generate RDI-L within 125 μ s of detecting an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2 and Figures 6-4 through 6-13). The LTE shall generate RDI-L by inserting the code '110' in bits 6, 7, and 8 of the K2 byte.

R6-201 [550v2] If bits 6 through 8 of the K2 byte are not used for other purposes (e.g., the linear APS mode indication), the LTE shall deactivate RDI-L by inserting the code '000' in bits 6, 7, and 8 of the K2 byte within 125 μ s of terminating the defect that caused it to be sent (assuming it has been sent for any minimum RDI-L assertion time supported by the NE, see below).

R6-202 [551v2] If bits 6 through 8 of the K2 byte are used for other purposes, LTE shall deactivate RDI-L by inserting an "appropriate code" (see below) in those bits within 125 μ s of terminating the defect that caused it to be sent (assuming it has been sent for the minimum assertion time).

For LTE that uses bits 6 through 8 of the K2 byte to indicate its linear APS mode, the "appropriate code" referred to in **R6-202 [551v2]** can be either the mode indication, or the code '000' for 50 to 200 ms, followed by the mode indication. In some (but not all) situations, inserting the code '000' for 50 to 200 ms will allow new LTE to interoperate with older LTE that was built to early issues of Bellcore's SONET criteria documents. (In those documents, an RDI-L defect was required to be terminated only if the '000' code was received.)

Note that the RDI-L generation and deactivation times in **R6-200 [549v2]**, **R6-201 [550v2]**, and **R6-202 [551v2]** are more stringent than the RDI-L specifications that appear in ANSI T1.105 (which states only that RDI-L must be generated or removed within 100 ms of detecting or terminating an incoming defect). However, it is not expected that these requirements will be changed. Three reasons for this are:

- The 125 μ s RDI-L generation and removal times in these requirements are consistent with the times that appeared in previous issues of this document
- Many existing implementations conform to these time requirements
- LTE that conforms to these requirements can also meet the standard (if it also conforms to **O6-203 [956]**).

With the definition and increasing support of far-end Line PM parameters whose accumulation is affected by the presence or absence of an RDI-L defect (see Section 6.2.2.4.2), it is important for LTE that generates an RDI-L signal to continue to generate it long enough to assure that its peer LTE will detect an RDI-L defect. ANSI T1.105 specifies that LTE must detect an RDI-L defect if the '110' code is received in 10 consecutive frames, and therefore specifies an RDI-L minimum assertion time of 20 frames (i.e., the defect will still be detected even if transmission errors occur that affect the '110' code in any one frame). To better align with T1.105 in this area, the following criteria have been added (**O6-203 [956]**) and modified (**R6-204 [552v2]** and **R6-205 [553v2]**).

O6-203 [956] When LTE generates RDI-L, it should generate it for at least 20 frames.

- R6-204** [552v2] LTE shall detect an RDI-L defect when bits 6, 7, and 8 of the K2 byte contain the '110' pattern in five to ten consecutive frames.
- R6-205** [553v2] LTE shall terminate the RDI-L defect when bits 6, 7, and 8 of the K2 byte contain any pattern other than the code '110' in five to ten consecutive frames.

To test conformance to **R6-204 [552v2]** and **R6-205 [553v2]**, it is sufficient to apply a consistent non-'110' code.

- R6-206** [554] An NE shall declare an RFI-L failure when an RDI-L defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-L failure, the NE shall set an RFI-L failure indication for the line and (if RFI-L is a Reported failure, see Section 6.2.1.8) send a message to an OS.
- R6-207** [555] An NE shall clear the RFI-L failure when the RDI-L defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the RFI-L failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.3.2 STS Path Remote Defect Indication (RDI-P) and Remote Failure Indications (RFI-P)

An RDI-P signal indicates to STS PTE that its peer STS PTE has detected a defect on the signal that it (the first STS PTE) originated. Three bits are reserved for the RDI-P signal (as shown in Table 6-4), and detected RDI-P defects are used to derive RFI-P failures.

The details of the signal that STS PTE uses to notify upstream STS PTE that it has detected a defect have undergone a number of changes as the SONET standards and Bellcore criteria have evolved (i.e., from an "STS Path Yellow" signal using bit 5 of the G1 byte, to a one-bit RDI-P signal using bit 5 of G1, to the enhanced two-bit RDI-P signal using bits 5 and 6 of G1 that appeared in Issue 1 of this document). Currently an enhanced, three-bit RDI-P signal using bits 5 through 7 of G1 is defined in ANSI T1.105. The specifications in that standard are not expected to change, and therefore this section contains criteria that are relevant to that enhanced RDI-P signal. However, criteria for a one-bit RDI-P signal using bit 5 of G1 are included to cover products that have not implemented the enhanced version of RDI-P. Note that in this document, "RDI-P" is used generically to refer to both the one-bit and enhanced versions in text and criteria that are not specific to one particular version. In cases where it is necessary to distinguish between the two versions, the one-bit version is referred to as "one-bit RDI-P", while the enhanced version is referred to as "ERDI-P". Similarly, "RFI-P" is used to generically refer to a failure derived from a (generic) RDI-P defect, while "one-bit RFI-P" and "ERFI-P" are used in cases that are specific to a particular version.

Although ERDI-P was originally intended to be "backward compatible" with one-bit RDI-P, in the final version there are actually some significant differences. These differences are expected to primarily affect the accumulation of near-end and far-end STS Path PM. Therefore in applications where the STS Path PM data is expected to be of importance, it is highly recommended that the STS PTE at both ends of the path (and any intermediate NEs provisioned to perform intermediate-path PM on that path) support the same version of RDI-P. In addition, it is recommended that the supported version be ERDI-P.

O6-208 [957v2] An NE should support ERDI-P generation and detection.

R6-209 [556v2] STS PTE shall generate an appropriate RDI-P signal, as specified in Table 6-4, within 100 ms of detecting a listed defect. (Also see Figures 6-5 through 6-13.)

O6-210 [557v2] STS PTE should generate an appropriate RDI-P signal as specified in Table 6-4 within 125 μ s of detecting a listed defect.

Note that for the ERDI-P codes shown in Table 6-4, bits 6 and 7 of the G1 byte are always set to opposite values. Conversely, in the one-bit RDI-P signal those two bits must be set to the same value (and should be set to '00', see Section 3.2). With RDI-P defined as described, STS PTE that supports ERDI-P can determine which version of RDI-P the far-end STS PTE is sending from the incoming G1 bits 6 and 7. If the received bits 6 and 7 are set to either '01' or '10', then the far-end STS PTE is sending ERDI-P and the near-end monitors G1 bits 5 through 7 to detect and terminate RDI-P defects. If the received bits 6 and 7 are set to either '00' or '11', then the far-end STS PTE is sending one-bit RDI-P and the near-end monitors only G1 bit 5 to detect and terminate RDI-P defects.

R6-211 [958] If ERDI-P is supported and the STS PTE has detected two or more of the listed defects, it shall generate the higher priority ERDI-P code based on the priorities shown in Table 6-4.

For example, if STS PTE that supports ERDI-P has detected an UNEQ-P defect (and is therefore sending the '110' code upstream) and then it detects an AIS-P defect, it must change its transmitted ERDI-P code to '101' within the 100 ms time period required by **R6-209** [556v2]. (Also see **O6-210** [557v2], **R6-212** [558] and **O6-213** [959].)

R6-212 [558] When STS PTE generates a particular type of RDI-P signal, it shall generate it for at least 10 frames.

O6-213 [959] When STS PTE generates a particular type of RDI-P signal, it should generate it for at least 20 frames.

Note that the applicability of these criteria (i.e., **R6-212** [558] and **O6-213** [959]) is currently under study for the case where a higher priority incoming defect is detected before

the ERDI-P code for a lower priority defect has been sent for the specified time (see GR-253-ILR Issue ID 253-140).

- R6-214** [559v2] STS PTE shall deactivate (or change, as appropriate) an RDI-P signal within 100 ms of terminating the defect that caused it to be generated.
- O6-215** [560v2] STS PTE should deactivate the RDI-P signal within 125 μ s of terminating the defect that caused it to be sent (assuming that the signal has been sent for the minimum assertion time supported by the NE, see **R6-212** [558] and **O6-213** [959]).
- R6-216** [561] If **O6-210** [557v2] and **O6-215** [560v2] are not both met, then the delay time to generate an RDI-P signal (i.e., the time between detection of the defect and generation of the RDI-P signal) shall be within 500 μ s of the delay time to deactivate the RDI-P signal (i.e., the time between termination of the defect and deactivation of the RDI-P signal).
- R6-217** [960] STS PTE that does not support ERDI-P shall detect a one-bit RDI-P defect when a '1' is received in bit 5 of G1 for ten consecutive frames.
- R6-218** [562v2] STS PTE that supports ERDI-P shall detect an RDI-P defect when one of the "RDI-P defect" codes shown in Table 6-4 (one-bit or enhanced) is received for five to ten consecutive frames.
- R6-219** [961] STS PTE that does not support ERDI-P shall terminate the one-bit RDI-P defect when a '0' is received in bit 5 of G1 for ten consecutive frames.
- R6-220** [564v2] STS PTE that supports ERDI-P shall terminate a particular type of RDI-P defect (one-bit or enhanced) when a code other than the code corresponding to that defect is received for five to ten consecutive frames.
- Note that in some situations, STS PTE that supports ERDI-P may simultaneously terminate one RDI-P defect and detect a different RDI-P defect.
- R6-221** [566v2] An NE shall declare the corresponding one-bit RFI-P or ERFI-P failure when a particular type of RDI-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-P failure, the NE shall set an RFI-P failure indication for the STS path and (if RFI-P is a Reported failure, see Section 6.2.1.8) send a message to an OS unless the condition in **R6-286** [629v3] applies.

- R6-222** [962v2] If ERDI-P is supported, the RFI-P failure indication shall indicate if the failure was derived from an incoming Server, Connectivity or Payload ERDI-P defect, or a one-bit RDI-P defect (see Table 6-4).
- R6-223** [567v2] An NE shall clear the corresponding RFI-P failure when the particular type of RDI-P defect that caused it to be declared is absent for 10 (± 0.5) seconds. Upon clearing the RFI-P failure, the NE shall clear the RFI-P failure indication and send a clear message to the OS (if the failure was reported to the OS).

Table 6-4. RDI-P Bit Settings and Interpretation

G1 Bits 5, 6 and 7	Priority of ERDI-P Codes	Trigger	Interpretation
0xx ^a	Not Applicable	No defects	No RDI-P defect
1xx ^a	Not Applicable	AIS-P ^b , LOP-P	one-bit RDI-P defect ^c
001 ^d	4	No defects	No RDI-P defect
010 ^d	3	PLM-P, LCD-P ^e	ERDI-P Payload defect
101 ^d	1	AIS-P ^b , LOP-P	ERDI-P Server defect
110 ^d	2	UNEQ-P, TIM-P ^f	ERDI-P Connectivity defect

Notes:

- This code is transmitted by STS PTE that does not support ERDI-P. If ERDI-P is not supported, G1 bits 6 and 7 must be set to the same value, and should be set to '00' (see Section 3.2).
- In the functional model used in this document, incoming defects that are detected at STE or LTE (i.e., LOS, LOF, AIS-L) cause AIS-P to be generated downstream. This AIS-P is then detected by the STS PTE and causes it generate RDI-P.
- The particular type of defect that should be detected by STS PTE that supports ERDI-P when it receives these codes is currently under study (see GR-253-ILR Issue ID 253-120).
- This code is transmitted by STS PTE that supports ERDI-P.
- Loss of Cell Delineation (LCD-P) is a Payload defect that is defined in ANSI T1.646-1995, *Broadband ISDN - Physical Layer Specification for User-Network Interfaces Including DS1/ATM*, for ATM payloads carried in SONET signals using the ATM mappings shown in Sections 3.4.2.2, 3.4.3.1.3, and 3.4.3.2.1. The criteria that indicate that ERDI-P is to be generated when an LCD-P defect is detected are currently under study. Further information on the LCD-P defect may be included in a future issue of this document.
- If TIM-P detection is activated for the terminating path. Also see GR-253-ILR Issue ID 253-139.

6.2.1.3.3 VT Path Remote Defect Indication (RDI-V) and Remote Failure Indication (RFI-V)

An RDI-V signal is used in all VT-based applications to indicate to VT PTE that its peer VT PTE has detected a defect on the signal that it (the first VT PTE) originated. Four bits

are reserved for the RDI-V signal (as shown in Table 6-5), and in most applications (i.e., applications other than those using the byte-synchronous DS1 mapping) detected RDI-V defects are used to derive RFI-V failures.

The details of the signals that VT PTE uses to notify upstream VT PTE that it has detected a defect have undergone a number of changes as the SONET standards and Bellcore criteria have evolved (i.e., from an "VT Path Yellow" signal using bit 8 of the V5 byte, to a one-bit RDI-V signal using bit 8 of V5 and a one-bit RFI-V signal using bit 4 of V5, to an enhanced two-bit RDI-V signal using bits 4 and 8 of V5 and a one-bit RFI-V signal using bit 8 of the Z7 byte). Currently an enhanced, four-bit RDI-V signal using bit 8 of V5 and bits 5 through 7 of Z7, and an RFI-V signal using bit 4 of V5 are defined in ANSI T1.105. The specifications in that standard are not expected to change, and therefore this section contains criteria that are relevant to those signals. However, criteria for a one-bit RDI-V signal using bit 8 of V5 are included to cover products that have not yet implemented the enhanced version of RDI-V. Note that in this document, "RDI-V" is used to refer generically to both the one-bit and enhanced versions in text and criteria that are not specific to one particular version. In cases where it is necessary to distinguish between the two versions, the one-bit version is referred to as "one-bit RDI-V", while the enhanced version is referred to as "ERDI-V". Similarly, "RFI-V" is used to generically refer to a failure derived from a (generic) RDI-V defect, while "one-bit RFI-V" and "ERFI-V" are used in cases that are specific the particular version.

Although ERDI-V was originally intended to be "backward compatible" with one-bit RDI-V, in the final version there are actually some significant differences. These differences are expected to primarily affect the accumulation of near-end and far-end VT Path PM. Therefore in applications where the VT Path PM data is expected to be of importance, it is highly recommended that the VT PTE at both ends of the path (and any intermediate NEs provisioned to perform intermediate-path PM on that path) support the same version of RDI-V. In addition, it is recommended that the supported version be ERDI-V.

- O6-224** **[963v2]** An NE should support enhanced RDI-V generation and detection.
- R6-225** **[568v2]** VT PTE shall generate an appropriate RDI-V signal, as specified in Table 6-5, within 100 ms of detecting a listed defect. (Also see Figures 6-7 through 6-13.)
- O6-226** **[569v2]** VT PTE should generate an appropriate RDI-V signal as specified in Table 6-5 within 500 μ s of detecting a listed defect.
- R6-227** **[964]** If ERDI-V is supported and the VT PTE has detected two or more of the listed defects, it shall generate the higher priority ERDI-V code based on the priorities shown in Table 6-5.

For example, if VT PTE that supports ERDI-V has detected an UNEQ-V defect (and is therefore sending the '1' and '110' codes upstream) and then it detects an AIS-V defect, it must change its transmitted ERDI-V code to '1' and '101' within the 100 ms time period required by **R6-225 [568v2]**. (Also see **O6-226 [569v2]**, **R6-228 [570]** and **O6-229 [965]**.)

R6-228 [570] When VT PTE generates a particular type of RDI-V signal, it shall generate that signal for at least 10 superframes.

O6-229 [965] When VT PTE generates a particular type of RDI-V signal, it should generate it for at least 20 superframes.

Note that the applicability of these criteria (i.e., **R6-228 [570]** and **O6-229 [965]**) is currently under study for the case where a higher priority incoming defect is detected before the ERDI-V code for a lower priority defect has been sent for the specified time (see GR-253-ILR Issue ID 253-140).

R6-230 [571v2] VT PTE shall deactivate (or change, as appropriate) an RDI-V signal within 100 ms of terminating the defect that caused it to be sent.

O6-231 [572v2] VT PTE should deactivate the RDI-V signal within 500 μ s of terminating the defect that caused it to be sent (assuming that the signal has been sent for the minimum assertion time supported by the NE, see **R6-228 [570]** and **O6-229 [965]**).

R6-232 [573] If **O6-226 [569v2]** and **O6-231 [572v2]** are not both met, then the delay time to generate the RDI-V signal (i.e., the time between detection of the defect and generation of the RDI-V signal) shall be within 2 ms of the delay time to deactivate the RDI-V signal (i.e., the time between termination of the defect and deactivation of the RDI-V signal).

Note that for the ERDI-V codes shown in Table 6-5, bits 6 and 7 of the Z7 byte are always set to opposite values. Conversely, in the one-bit RDI-V signal those two bits must be set to the same value (and should be set to '00'). With RDI-V defined as described, VT PTE that supports ERDI-V can determine which version of RDI-V the far-end VT PTE is sending from the incoming Z7 bits 6 and 7. If the received bits 6 and 7 are set to either '01' or '10', then the far-end VT PTE is sending ERDI-V and the near-end monitors the Z7 bits to detect and terminate RDI-V defects. If the received bits 6 and 7 are set to either '00' or '11', then the far-end VT PTE is sending one-bit RDI-V and the near-end monitors V5 bit 8 to detect and terminate RDI-V defects.

R6-233 [966] VT PTE that does not support ERDI-V shall detect a one-bit RDI-V defect when a '1' is received in bit 8 of V5 for ten consecutive superframes.

- R6-234** [574v2] VT PTE that supports ERDI-V shall detect an RDI-V defect when one of the "RDI-V defect" codes shown in Table 6-5 (one-bit or enhanced) is received for five to ten consecutive VT superframes.
- R6-235** [967] VT PTE that does not support ERDI-V shall terminate the one-bit RDI-V defect when a '0' is received in bit 8 of V5 for ten consecutive superframes.
- R6-236** [576v2] VT PTE that supports ERDI-V shall terminate a particular type of RDI-V defect (one-bit or enhanced) when a code other than the code corresponding to that defect is received for five to ten consecutive superframes.

Note that in some situations, VT PTE that supports ERDI-V may simultaneously terminate one RDI-V defect and detect a different RDI-V defect.

- R6-237** [578v2] An NE that is not using the byte-synchronous DS1 mapping for a particular VT path shall declare the corresponding one-bit RFI-V or ERFI-V failure when a particular type of RDI-V defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-V failure, the NE shall set an RFI-V failure indication for the VT path and (if RFI-V is a Reported failure, see Section 6.2.1.8) send a message to an OS unless the condition in **R6-286** [629v3] applies.
- R6-238** [968v2] If ERDI-V is supported, the RFI-V failure indication shall indicate if the failure was derived from an incoming Server, Connectivity or Payload ERDI-V defect, or a one-bit RDI-V defect (see Table 6-5).
- R6-239** [579v2] An NE that is not using the byte-synchronous DS1 mapping for a particular VT path shall clear the corresponding RFI-V failure when the particular type of RDI-V defect that caused it to be declared is absent for 10 (± 0.5) seconds. Upon clearing the RFI-V failure, the NE shall clear the RFI-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

Table 6-5. RDI-V Bit Settings and Interpretation

Z7 Bits 5, 6 and 7	V5 Bit 8	Priority of ERDI-V Codes	Trigger	Interpretation
yxx ^a	0	Not Applicable	No defects	No RDI-V defect
yxx ^a	1	Not Applicable	AIS-V ^b , LOP-V	one-bit RDI-V defect ^c
001 ^d	0 ^e	4	No defects	No RDI-V defect
010 ^d	0 ^e	3	PLM-V	ERDI-V Payload defect
101 ^d	1 ^e	1	AIS-V ^b , LOP-V	ERDI-V Server defect
110 ^d	1 ^e	2	UNEQ-V, TIM-V ^f	ERDI-V Connectivity defect

Notes:

- This code is transmitted by VT PTE that does not support ERDI-V. If ERDI-V is not supported, Z7 bits 6 and 7 must be set to the same value, and should be set to '00' (see Section 3.2).
- In the functional model used in this document, incoming defects that are detected at STE, LTE or STS PTE (e.g., LOS, LOF, AIS-L, AIS-P) cause AIS-V to be generated downstream. This AIS-V is then detected by the VT PTE and causes it generate RDI-V.
- The particular type of defect that should be detected by VT PTE that supports ERDI-V when it receives these codes is currently under study (see GR-253-ILR Issue ID 253-120).
- This code is transmitted by VT PTE that supports ERDI-V.
- V5 bit 8 is set to the same value as Z7 bit 5 by the VT PTE that supports ERDI-V. At the receiving VT PTE, V5 bit 8 is ignored unless Z7 bits 6 and 7 are both set to '0' or both set to '1'.
- Although a definition of TIM-V currently does not appear in either this document or ANSI T1.231-1997, one could be included in a future issue of this document. See GR-253-ILR, Issue ID 253-94 for a discussion of the effect that a TIM-V defect would have on various features supported by SONET NEs.

In addition to RDI-V signals, an RFI-V signal is also defined for use in applications where the byte-synchronous DS1 mapping is used. In those applications, it is the reception of the RFI-V signal that causes the declaration of the RFI-V failure for the path. This declaration can then be used to initiate trunk conditioning when required. Also, in some applications RFI-V signals are translated to and from DS1 RAI signals at a DS1 interface. Bit 4 of the V5 byte now carries the RFI-V signal.

R6-240 [580v3] VT PTE that is using the byte-synchronous DS1 mapping for a particular path shall generate an RFI-V signal by setting bit 4 of the V5 byte to '1' within 500 μ s of declaring an AIS-V failure (or a lower-layer, traffic-related, near-end failure, see Section 6.2.1.8.2), an LOP-V failure, an UNEQ-V failure, or a PLM-V failure (as shown in Figures 6-9 through 6-13).

R6-241 [969] VT PTE that is byte-synchronously mapping a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall generate an RFI-V signal by

setting bit 4 of the V5 byte to '1' within 500 μ s of declaring a DS1 RAI failure for an incoming DS1 (as shown in Figure 6-10).

- R6-242** [581v2] VT PTE that is using the byte-synchronous DS1 mapping shall deactivate the RFI-V signal by setting bit 4 of V5 to '0' within 500 μ s of clearing the failure that caused the RFI-V signal to be sent.
- R6-243** [582v2] An NE that is using the byte-synchronous DS1 mapping for a particular VT path shall declare an RFI-V failure upon receiving a '1' in bit 4 of V5 for 10 consecutive superframes (i.e., for 5 ms). Upon declaring the RFI-V failure, the NE shall set an RFI-V failure indication for the VT path and (if RFI-V is a Reported failure, see Section 6.2.1.8) send a message to an OS unless the condition in **R6-286** [629v3] applies.
- R6-244** [583v2] An NE shall clear the RFI-V failure within 50 ms of receiving a '0' in bit 4 of V5 for 10 consecutive superframes. Upon clearing the RFI-V failure, the NE shall clear the RFI-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

For testing an NE's conformance to **R6-244** [583v2], it is sufficient to send '0' in bit 4 of V5 for less than 10 consecutive superframes (the RFI-V failure must not be cleared), and for more than 50 ms (the RFI-V failure must be cleared).

Finally, VT PTE that uses the byte-synchronous DS1 mapping for a particular VT path may optionally provide the ability to use the RDI-V signal to derive RFI-V failures, in addition to the requirement to use the RFI-V signal. However, only the RFI-V signal is used to initiate trunk conditioning (see Section 6.2.1.6).

6.2.1.3.4 DS_n RDI and RAI Signals

In some applications a SONET NE may need to generate or detect DS_n RDI or RAI signals. Figures 6-8 through 6-13 (in Section 6.2.1.7) illustrate the use of DS_n RDI and RAI for STS and VT PTE terminating SPEs of various compositions. GR-499-CORE contains the requirements on the construction of DS1, DS1C, DS2 and DS3 RAI, and DS3 RDI.⁸ DS0 RAI is defined in SONET as a specific pattern in the ABCD signaling bits (i.e., A = 0, B = 1, C = 1, D = 1). Note that DS0 RAI may not be applicable outside of SONET networks, so an interface to a non-SONET NE may require the application of service-specific trunk conditioning codes.

8. DS3 RDI uses the X-bits in the DS3 frame, and is defined for M23 and C-bit parity applications. DS3 RAI is defined only in C-bit parity applications, and is also referred to as Far End Alarm and Control (FEAC) signals. See GR-499-CORE for additional information.

- CR6-245** [970] A SONET NE that terminates an M23 application DS3 path may be required to generate DS3 RDI upstream upon detecting certain defects, as shown in Figure 6-8.
- R6-246** [971] A SONET NE that terminates a C-bit parity application DS3 path shall generate DS3 RDI upstream upon detecting certain defects, as shown in Figure 6-8.
- R6-247** [972] A SONET NE shall remove a DS3 RDI upon terminating the defect that caused it to be generated.
- R6-248** [584v2] A SONET NE that terminates a DS_n path other than an M23 application DS3 path shall generate DS_n RAI upstream (unless it is provisioned to apply a service-specific trunk conditioning code upstream) upon declaring certain failures, as shown in Figures 6-8, 6-12 and 6-13.
- R6-249** [586v2] VT PTE provisioned to byte-synchronously map a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall generate DS1 RAI downstream upon declaring an RFI-V failure as a result of receiving an RFI-V signal, as shown in Figure 6-9.

As discussed in Section 6.2.1.3.3, VT PTE's generation of DS1 RAI is based on the detection of an incoming RFI-V signal and the subsequent declaration of an RFI-V failure, not on RFI-V failures derived from detected RDI-V defects.

- R6-250** [973] A SONET NE that provides access to and processing of individual DS0 channels shall generate DS0 RAI upstream upon declaring certain failures that cause trunk conditioning to be applied downstream (unless it is provisioned to also apply a service-specific trunk conditioning code upstream), as shown in Figures 6-11 and 6-12.

Note that DS0 RAI is not sent upstream if DS0 AIS is generated (or allowed to pass) downstream, or if the NE is also provisioned to apply trunk conditioning in the upstream direction. It is sent upstream only if the NE is provisioned to apply trunk conditioning only in the downstream direction.

- R6-251** [974] A SONET NE that provides access to and processing of individual DS0 channels shall generate DS0 RAI downstream when it declares a DS1 RAI failure (unless it is provisioned to apply a service-specific trunk conditioning code), as shown in Figure 6-12.
- R6-252** [975] A SONET NE shall remove a DS_n RAI upon clearing the failure that caused it to be generated.

In addition to generating and removing DS_n RDI and RAI, in some applications a SONET NE may also need to be able to monitor for DS_n RDI and RAI (e.g., to detect incoming DS_n RAI signals, and declare and clear DS_n RAI failures).

- CR6-253** [976] A SONET NE that terminates M23 application DS3 paths may be required to detect DS3 RDI for those paths.
- R6-254** [977] A SONET NE that terminates C-bit parity application DS3 paths shall detect DS3 RDI for those paths.
- R6-255** [978] A SONET NE that terminates DS_n paths shall detect DS_n RAI signals (if defined) for those paths.
- R6-256** [585v2] VT PTE that byte-synchronously maps a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall detect a DS1 RAI signal received at the DS1 interface (and insert RFI-V downstream as shown in Figure 6-10).
- R6-257** [979] A SONET NE that provides access to and processing of individual DS0 channels, and that is provisioned to apply trunk conditioning in a particular direction shall detect DS0 RAI signals in that direction (and apply the trunk conditioning code).

Note that if the NE is not provisioned to apply trunk conditioning in a particular direction, then DS0 RAI is simply passed through in that direction and does not need to be detected by the NE.

If an NE detects DS0 RAI signals, the following requirements apply.

- R6-258** [587v2] A SONET NE that detects DS0 RAI signals shall declare a DS0 RAI failure if it receives two consecutive sets of ABCD signaling bits set to the code '0111' (i.e., the code '0111' persisting for 6 ms).
- R6-259** [588v2] A SONET NE shall clear a DS0 RAI failure if it receives four consecutive sets of ABCD signaling bits set to any code other than the '0111' code (i.e., any code other than '0111' persisting for 12 ms).

6.2.1.4 Payload Defect Indication (PDI)

PDI is an application-specific code inserted by PTE to indicate to downstream equipment that there is a defect in one or more of its directly mapped embedded payloads. PDI is expected to be used primarily in ring interconnection applications, and possibly in some dual-homing applications. For example, in an application where DCSs are used to perform VT-level grooming between two dual-interconnected STS level UPSRs, the DCSs and all of the NEs on the rings would have to support PDI generation and/or detection. Note that

PDI is currently only defined at the STS Path level (i.e., PDI-P). Several options have been proposed for a VT level PDI (i.e., PDI-V carried as a code in the VT signal label, or as AIS-V); however, those options have not been accepted in ANSI and are not included here.

6.2.1.4.1 STS Payload Defect Indication (PDI-P)

PDI-P is a set of application-specific codes contained in the STS POH generated by STS PTE. It is used to indicate to downstream equipment (e.g., to the path selector in a downstream STS-level UPSR NE) that there is a defect in one or more of the directly mapped payloads contained in that STS SPE. The PDI-P codes appear in the STS Signal Label (C2 byte). Valid codes are shown in Table 3-3.

CR6-260 [591v2] STS PTE may be required to support PDI-P signal generation.

CR6-261 [589v2] STS PTE that supports PDI-P generation may be required to be provisionable as to whether it sends PDI-P signals.

The capability to disable PDI-P generation would be needed (for example) for backwards compatibility with older NEs that would consider any of the PDI-P codes to be mismatched to the locally provisioned STS PTE functionality.

R6-262 [980] If STS PTE that supports PDI-P generation and a VT-structured STS SPE does not process VT pointers, it shall (non-intrusively) detect and terminate LOP-V defects as if it were processing the those pointers (i.e., according to the criteria in Section 6.2.1.1.3).

R6-263 [981] STS PTE that supports PDI-P generation and a VT-structured STS SPE shall (non-intrusively) detect and terminate AIS-V defects as if it were VT PTE (i.e., according to the criteria in Section 6.2.1.2.3).

R6-264 [982] STS PTE that supports PDI-P generation and the asynchronous DS3 mapping shall (non-intrusively) detect and terminate DS3 AIS as if it were terminating the DS3 path (i.e., according to the criteria in Section 6.2.1.2.4).

CR6-265 [593v2] STS PTE that supports PDI-P generation and the asynchronous DS3 mapping may be required to be provisionable to (non-intrusively) detect and terminate DS3 OOF as if it were terminating the DS3 path (i.e., according to the criteria in Section 6.2.1.1.2).

See Section 6.2.1.8.7 for criteria concerning the declaration and clearing of failures based on the defects detected and terminated according to the preceding criteria.

- R6-266** [983] As a default STS PTE shall not monitor for DS3 OOF for the purposes of generating PDI-P.
- R6-267** [592v2] STS PTE that supports PDI-P generation shall generate (or change, as appropriate) the PDI-P signal within 100 ms of detecting an LOP-V, AIS-V, DS3 AIS, DS3 LOS, or (if so provisioned) DS3 OOF defect on any VT or DS3 payload that it embeds into the STS SPE that it is originating. The PDI-P signal shall be generated by inserting the code indicating the number of defective payloads as specified in Table 3-3.
- R6-268** [595v2] STS PTE that supports PDI-P generation shall deactivate (or change, as appropriate) the PDI-P signal within 100 ms of terminating a defect on one or more of its payloads.
- CR6-269** [984] A SONET NE may be required to support the detection of PDI-P signals.
- CR6-270** [985] A SONET NE that supports PDI-P detection may be required to be provisionable (on a per-STs path basis) as to whether it detects PDI-P.
- R6-271** [596v2] A SONET NE that supports the detection of PDI-P shall detect a PDI-P defect (or a change in the PDI-P defect, as appropriate) within 10 ms of the onset of at least five consecutive samples (which might not be consecutive frames) of STS Signal Labels (C2 bytes) containing a new PDI-P code.
- O6-272** [597v2] A SONET NE that supports the detection of PDI-P should detect a PDI-P defect (or a change in the PDI-P defect, as appropriate) immediately upon receipt of five consecutive frames of STS Signal Labels containing a new PDI-P code.
- R6-273** [598] A SONET NE that supports the detection of PDI-P shall terminate a PDI-P defect within 10 ms of the onset of at least five consecutive samples (which might not be consecutive frames) of STS Signal Labels that do not contain a PDI-P code.
- O6-274** [599] A SONET NE that supports the detection of PDI-P should terminate a PDI-P defect immediately upon receipt of five consecutive frames of STS Signal Labels that do not contain a PDI-P code.

Unlike most of the other defects defined in SONET, an NE that detects a persistent PDI-P defect is not required to declare a PDI-P failure or send a PDI-P failure report to an OS.⁹ The reason for this is that PDI-P defects are detected when there is a problem (i.e., a defect) with an embedded payload, not with the STS SPE itself. The root-cause defect will

normally be reported to the OS (if it persists so that a failure is declared) by some NE other than the one that *detects* the PDI-P defect. For example, an NE that detects a DS3 LOS, and *inserts* the appropriate PDI-P code (and DS3 AIS) downstream will normally declare a DS3 LOS failure.

Since PDI-P defects are expected to be used for STS path protection switching purposes in some applications, they can have a direct impact on the traffic carried by a SONET system and the user may need to determine the particular codes that are being transmitted and received. Section 6.2.3.2.3.B contains the criteria on diagnostics to report the STS Signal Label.

6.2.1.4.2 VT Payload Defect Indication (PDI-V)

PDI-V is currently not defined in SONET.

6.2.1.5 Maintenance Signals for Other Mappings

The criteria for DS4NA, FDDI, ATM, and DQDB maintenance signals are for further study.

6.2.1.6 Trunk Conditioning

Trunk conditioning is used during carrier failures to release connections, terminate customer billing, and remove trunks from service by instructing the switch or customer PBX to go idle and make the trunks busy. In a digital bit stream, trunk conditioning codes (i.e., one or more provisioned ABCD signaling codes and an 8-bit insertion word) are inserted into each DS0 channel affected by the carrier failure.¹⁰ For incoming signal defects that would otherwise cause DS0 AIS to be inserted or passed downstream, the signaling is frozen (i.e., the last values for the signaling bits are used) when the defect is initially detected. Then the defect is soaked for a time interval of approximately 2.5 seconds (or 3.25 seconds for DS0 AIS defects) to verify that a failure should be

-
9. If an NE does declare PDI-P failures, it is expected that the process used will be consistent with the failure declaration processes required for other defects/failures defined in this document. In addition, PDI-P defects are generally expected to be used for STS path protection switching purposes and the reader should see the appropriate application-specific GRs (e.g., GR-1230-CORE, GR-1400-CORE) for any additional criteria related to the reports an NE generates when it performs a path protection switch.
 10. The required bit pattern or patterns depend on the service being carried by the DS0 channel, and thus may differ for different services. Most services require a single (service-specific) code for trunk conditioning. Some require the application of one (service-specific) code followed by another to support a 2-stage (make-idle/make-busy) form of trunk conditioning. The trunk conditioning signaling codes are accompanied by an 8-bit insertion word. TR-NWT-000170 identifies signaling and 8-bit insertion codes for various types of access circuits.
-

declared. In addition, in most cases an integration timer is used to account for intermittent defects. For other signals that can trigger the application of trunk conditioning (i.e., RFI-V, DS1 RAI, and DS0 RAI signals), the detection of the signal causes a failure to be declared immediately (i.e., there is no soaking interval). In both cases, the trunk conditioning code or codes are applied to each affected DS0 when the failure is declared.

In SONET, trunk conditioning is applicable only when access to and processing of individual DS0 channels are provided in the SONET signal. Therefore, SONET NEs that support the byte-synchronous VT mapping in conjunction with DS0 rearrangement or termination functions need to support trunk conditioning functions. SONET NEs that support the byte-synchronous VT mapping but not DS0 rearrangement or termination functions do not need to support trunk conditioning functions. VT and DS1 level maintenance signaling is sufficient for those cases.

The maintenance signaling and trunk conditioning requirements related to byte-synchronously mapped DS1s at a SONET NE that performs DS0 rearrangement or DS0 termination are summarized in Figures 6-11 through 6-13. Note that Figures 6-9 and 6-10 apply to the same mapping at a DS1 interface where no DS0 rearrangement or termination functions are performed.

- R6-275** **[610v3]** An NE that provides access to and processing of individual DS0 channels shall set a red alarm when it declares an LOS, LOF, LOP-P, UNEQ-P, TIM-P (if activated), PLM-P, LOP-V, UNEQ-V, or PLM-V failure on the incoming signal.
- R6-276** **[612]** An NE shall clear a red alarm when the failure that caused it to be set has cleared.
- R6-277** **[986]** An NE that provides access to and processing of individual DS0 channels shall allow the user to provision it to apply trunk conditioning codes (for use in place of downstream DS0 AIS, downstream DS0 RAI, or upstream DS0 RAI, as applicable) on a per-DS0 basis. If the DS0 is not terminated, then the user shall be able to provision separate codes in each direction.

Note that trunk conditioning can only be applied in the upstream direction at a DS0 PTE.

- R6-278** **[609v2]** If it is provisioned to apply a trunk conditioning code on a particular DS0, the SONET NE shall freeze the DS0 signaling state upon detecting a defect that would otherwise cause DS0 AIS to be generated or passed downstream as shown in Figures 6-11 and 6-12.
- R6-279** **[611v2]** If a defect that causes a SONET NE to freeze DS0 signaling states persists so that the NE declares the associated failure, the NE shall apply the provisioned downstream trunk conditioning codes.

- R6-280** [987] If a defect that causes a SONET NE to freeze DS0 signaling states persists so that it declares a failure, the NE shall apply the provisioned upstream trunk conditioning codes (if any) instead of DS0 RAI.
- R6-281** [617v2] If it is provisioned to apply a trunk conditioning code on a particular DS0, the SONET NE shall apply the provisioned code downstream upon declaring an RFI-V, DS1 RAI, or DS0 RAI failure as a result of an incoming RFI-V, DS1 RAI, or DS0 RAI signal (as shown in Figures 6-11 and 6-12).
- R6-282** [618v2] A SONET NE shall remove trunk conditioning upon clearing the failure that had caused it to be applied.

6.2.1.7 Alarm-Related Events

Figures 6-3 through 6-14 illustrate the maintenance signals sent by SONET NEs when various defects are detected on incoming signals or certain equipment failures are declared. Those figures are followed by Figures 6-15 through 6-18, which illustrate the timing sequence for SONET NEs detecting defects, declaring failures, and setting their indications. These timing sequences are summarized below. Note that timing requirements regarding DS0 AIS and trunk conditioning-related signals and functions are not depicted in these particular figures (see Sections 6.2.1.2.4 and 6.2.1.6).

As Figures 6-15 through 6-18 indicate, downstream AIS and upstream RDI signals are sent immediately upon detecting the appropriate defect. Declaring a failure, setting the indication, and reporting these events to the OS are delayed for 2.5 (± 0.5) seconds to verify that the defect persists. Where the RFI-V signal is used, the signal is sent upstream when the failure is declared.

When a defect is terminated, downstream AIS and upstream RDI signals are removed immediately. When a failure is cleared on a signal using RFI-V, the upstream RFI-V signal is removed.

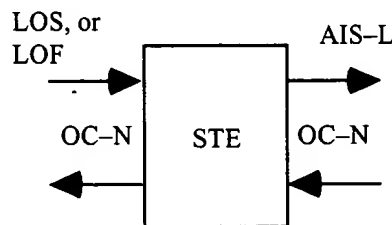
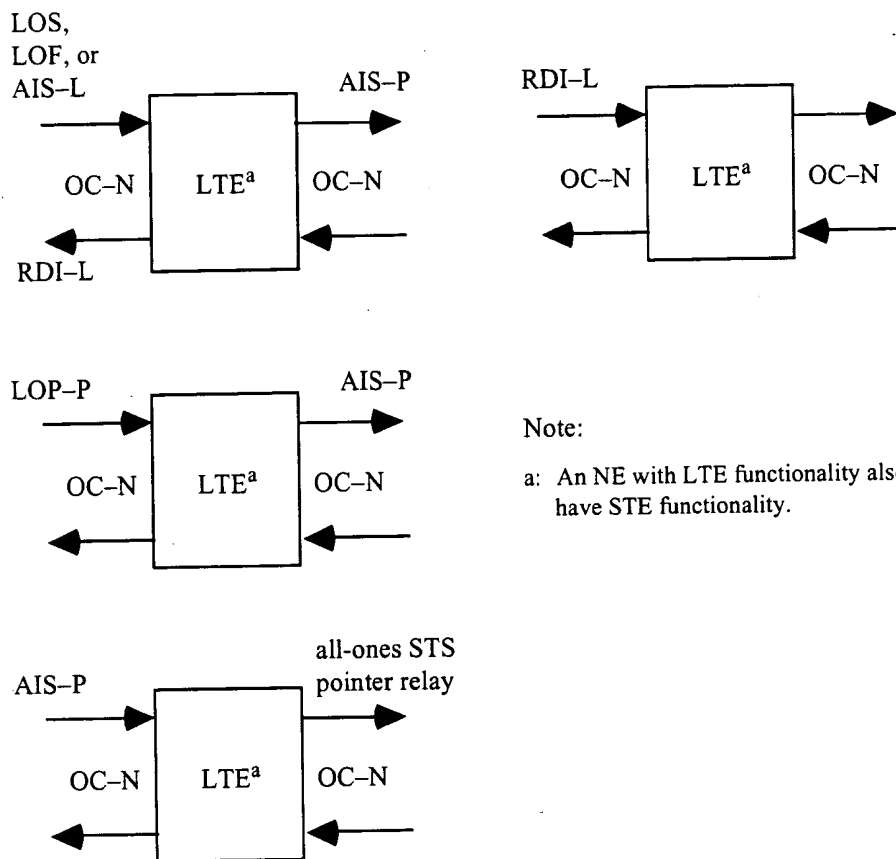


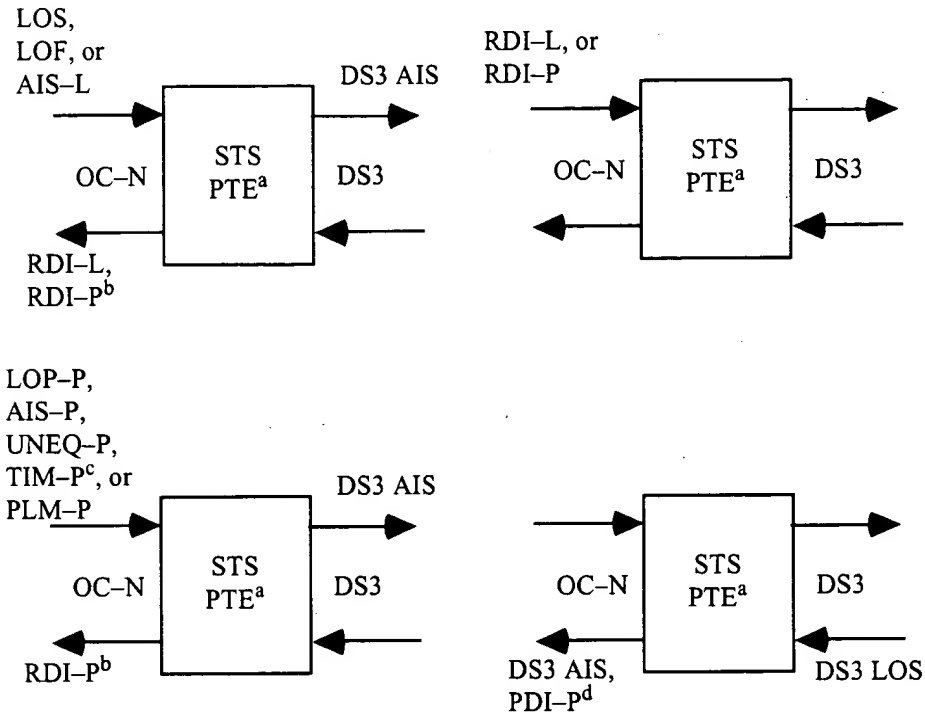
Figure 6-3. STE Maintenance Signals



Note:

a: An NE with LTE functionality also will have STE functionality.

Figure 6-4. LTE Maintenance Signals



Notes:

- a: An NE with STS PTE functionality also will have STE and LTE functionality.
- b: See Section 6.2.1.3.2 for the applicable RDI-P code for each type of incoming defect.
- c: If TIM-P detection is activated for the STS path. Also see GR-253-ILR Issue ID 253-139.
- d: Provisionable.
- e: Provisionable to cause generation of PDI-P and/or DS3 AIS.

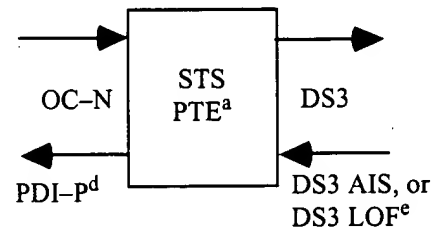
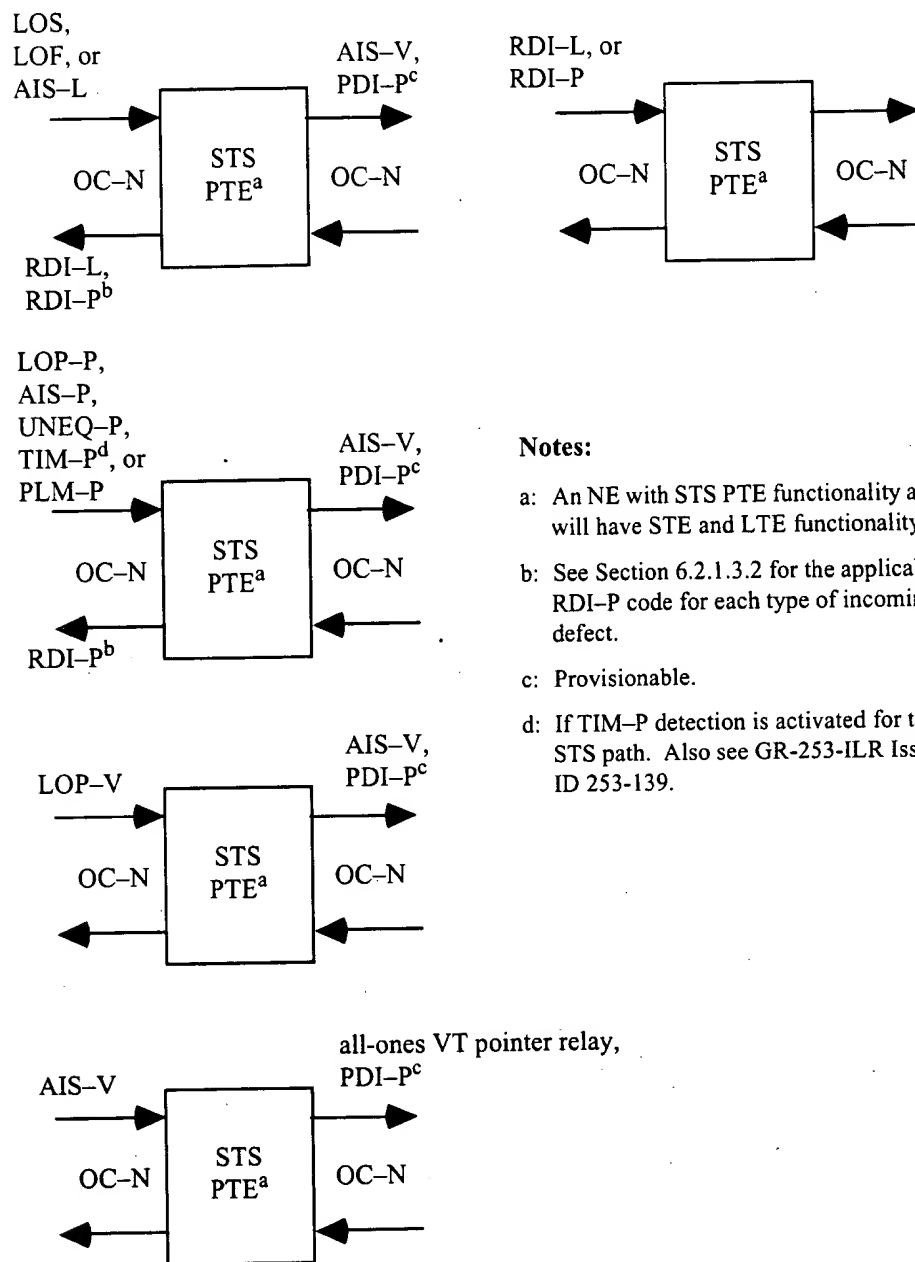


Figure 6-5. STS PTE (Asynchronous Mapping for DS3) Maintenance Signals



Notes:

- a: An NE with STS PTE functionality also will have STE and LTE functionality.
- b: See Section 6.2.1.3.2 for the applicable RDI-P code for each type of incoming defect.
- c: Provisionable.
- d: If TIM-P detection is activated for the STS path. Also see GR-253-ILR Issue ID 253-139.

Figure 6-6. STS PTE (VT-Structured STS-1 SPE) Maintenance Signals

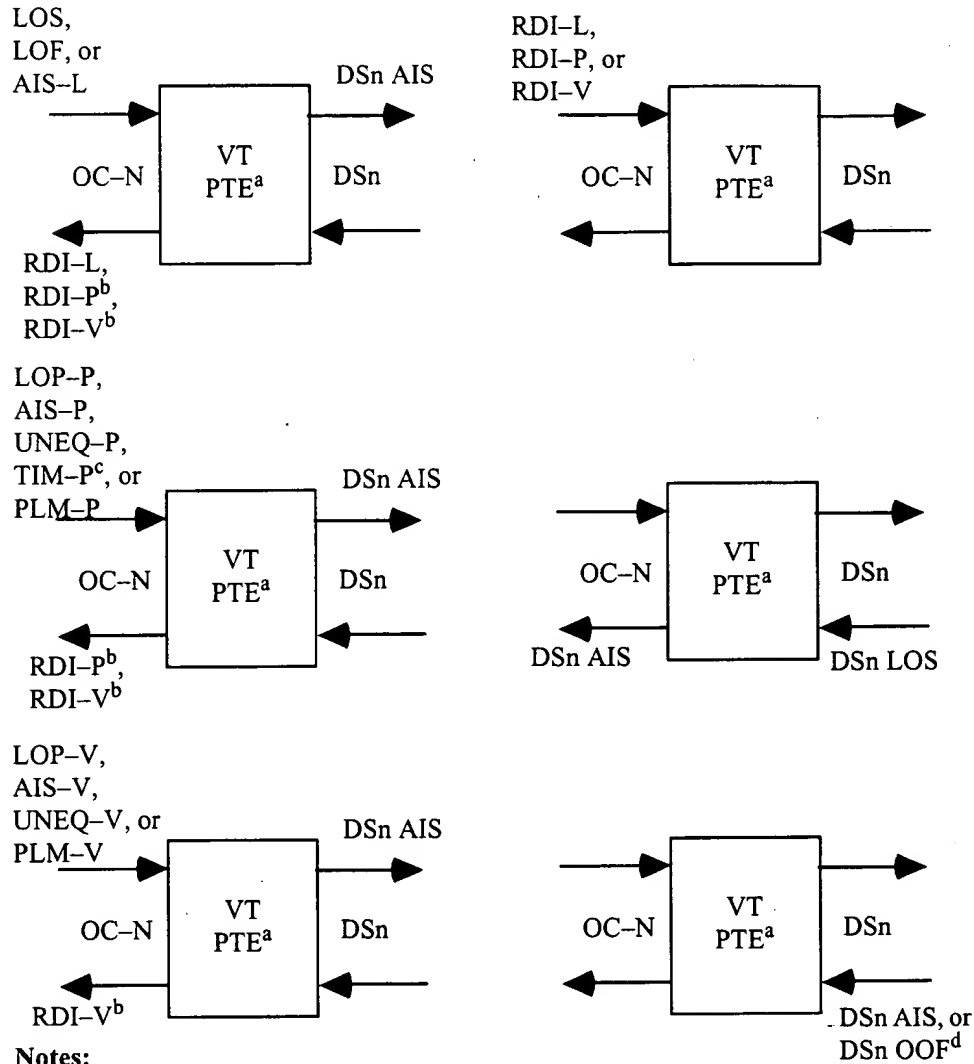
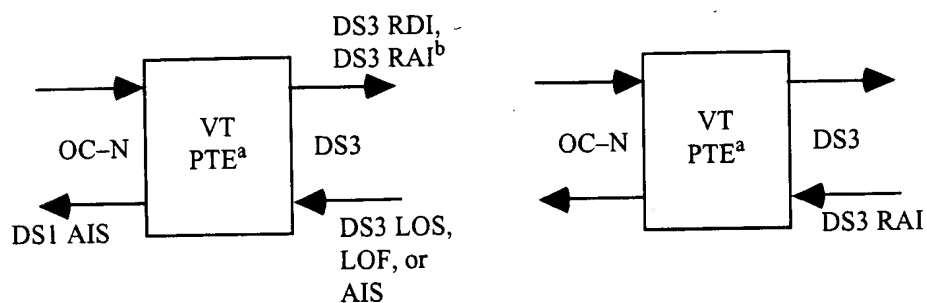


Figure 6-7. VT PTE (Asynchronous Mapping for DSn into VTx) Maintenance Signals, DSn Interface



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE, and STS PTE functionality.
- b: DS3 RAI is sent after timing the incoming defect.

Figure 6-8. VT PTE (Asynchronous Mapping for DS1 into VT1.5) Maintenance Signals, DS3 Interface with Embedded M13 Multiplex

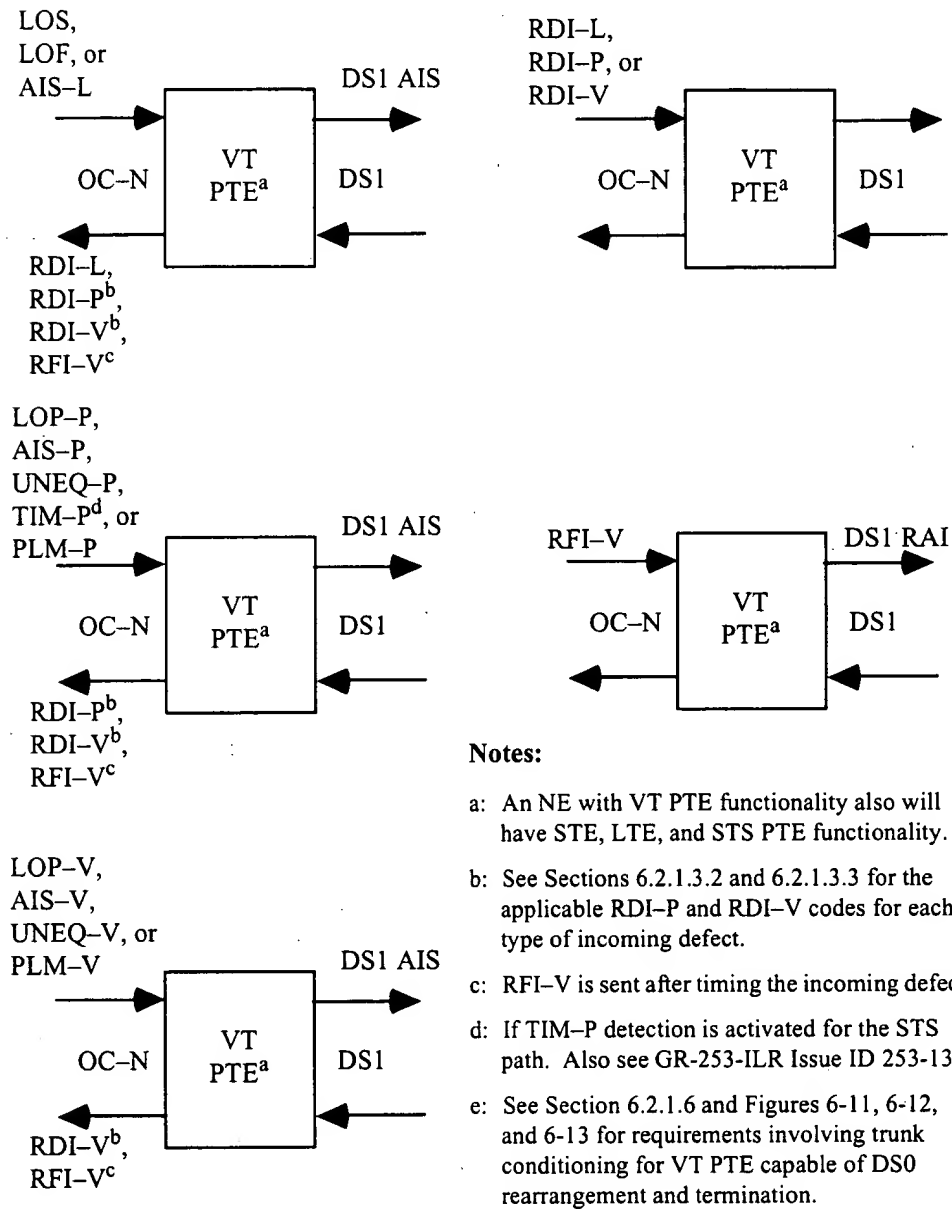
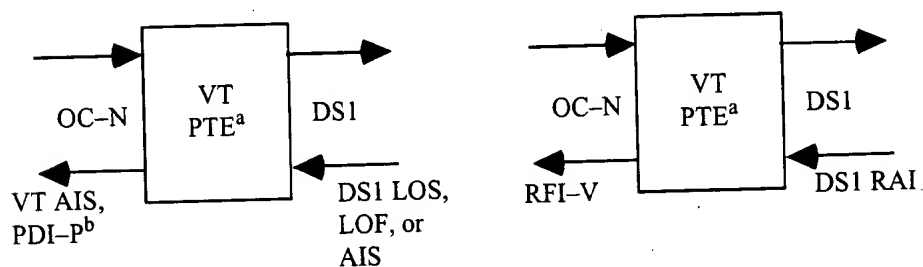


Figure 6-9. VT PTE (Byte-Synchronous Mapping for DS1 into VT1.5)
Maintenance Signals, DS1 Interface



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE, and STS PTE functionality.
- b: Provisionable.
- c: See Section 6.2.1.6 and Figures 6-11, 6-12, and 6-13 for requirements involving trunk conditioning for VT PTE capable of DS0 rearrangement and termination.

Figure 6-10. VT PTE (Byte-Synchronous Mapping for DS1 into VT1.5)
Maintenance Signals, DS1 Interface (Continued)

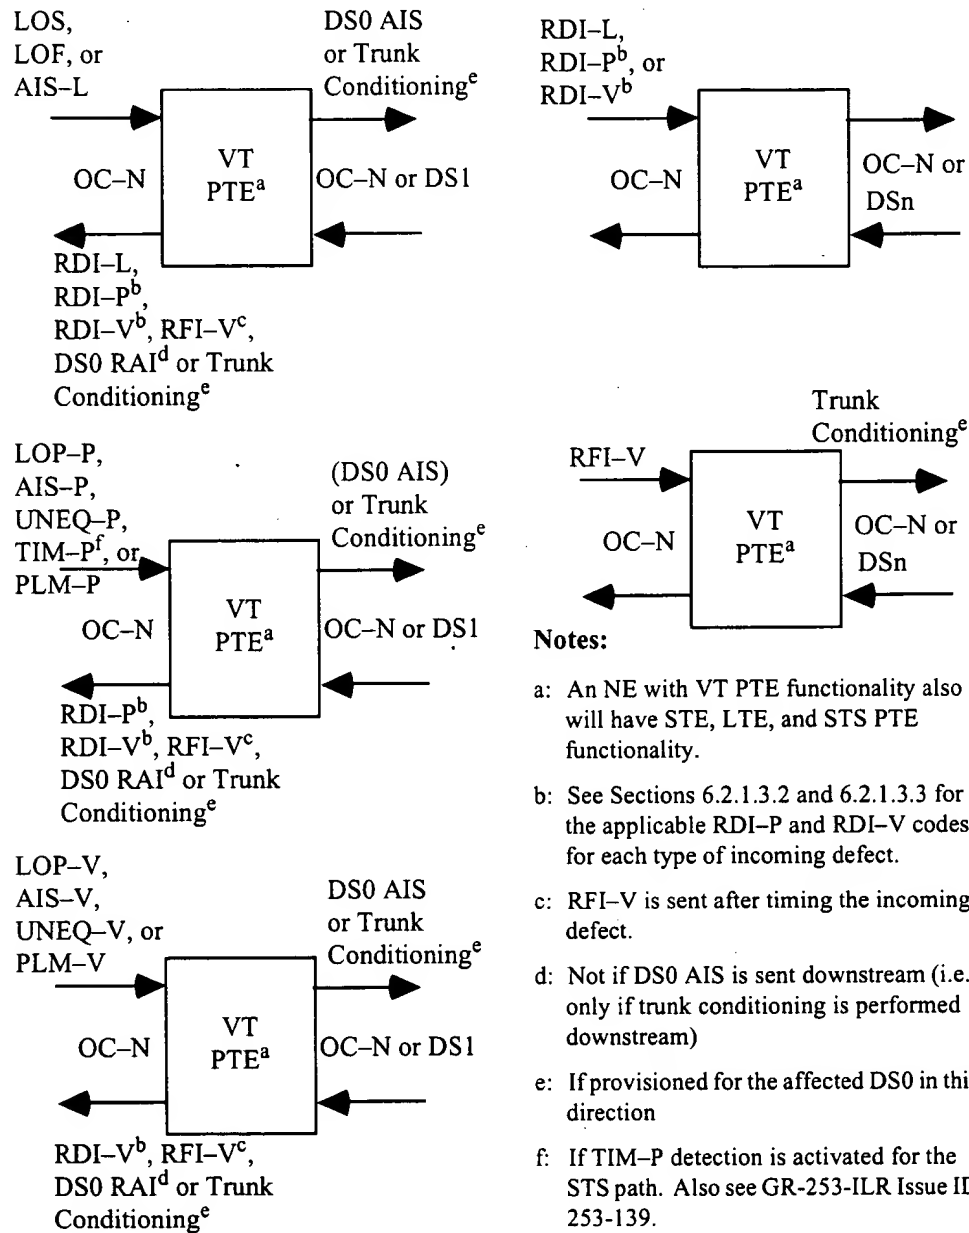
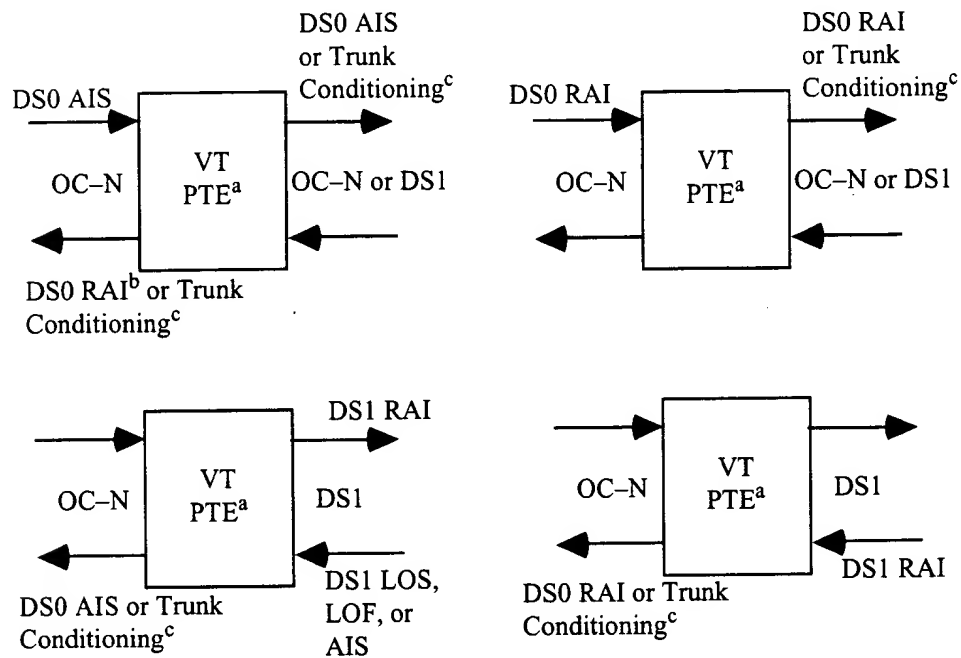


Figure 6-11. VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE, and STS PTE functionality.
- b: Not if DS0 AIS is sent downstream (i.e., only if trunk conditioning is performed downstream).
- c: If provisioned for the affected DS0 in this direction.

Figure 6-12. VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces (Continued)

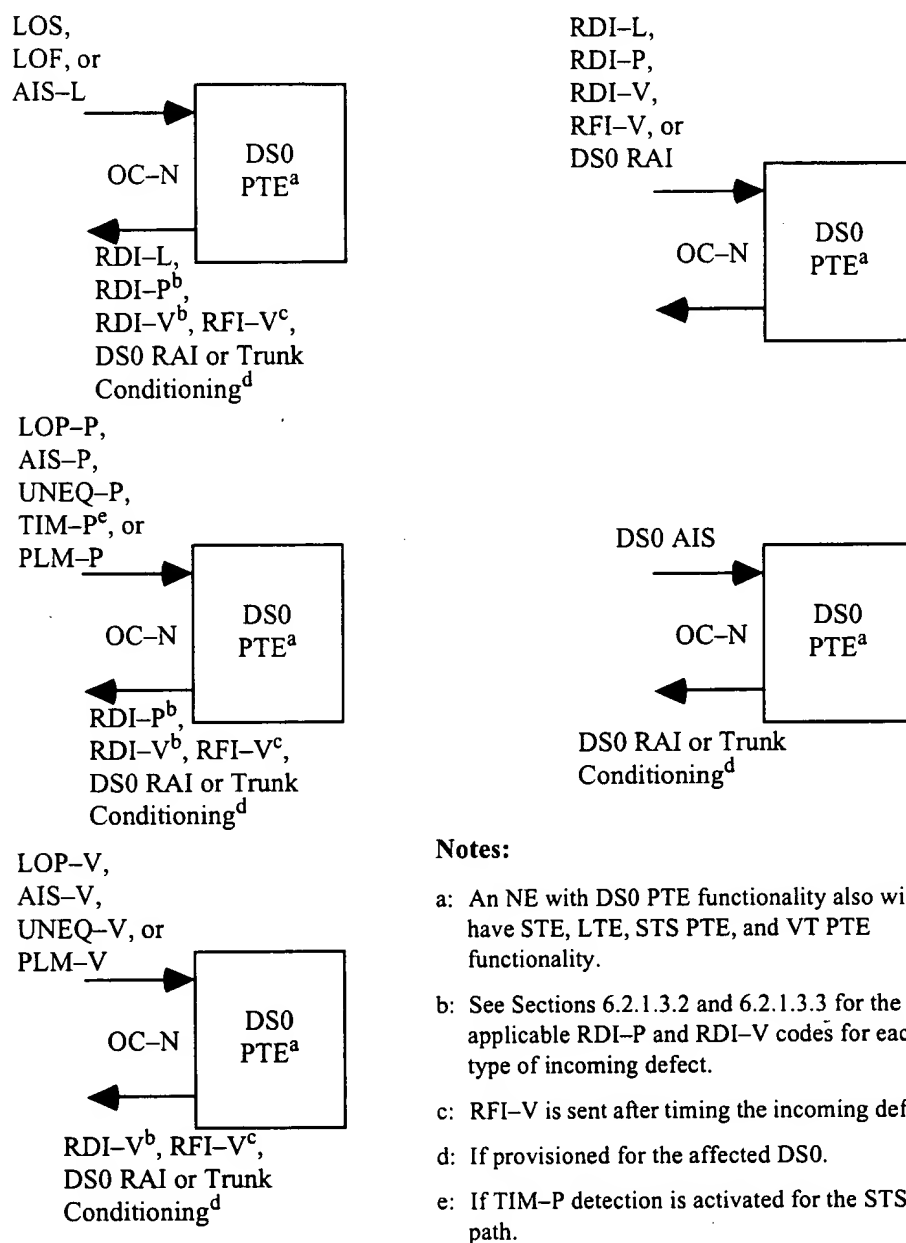
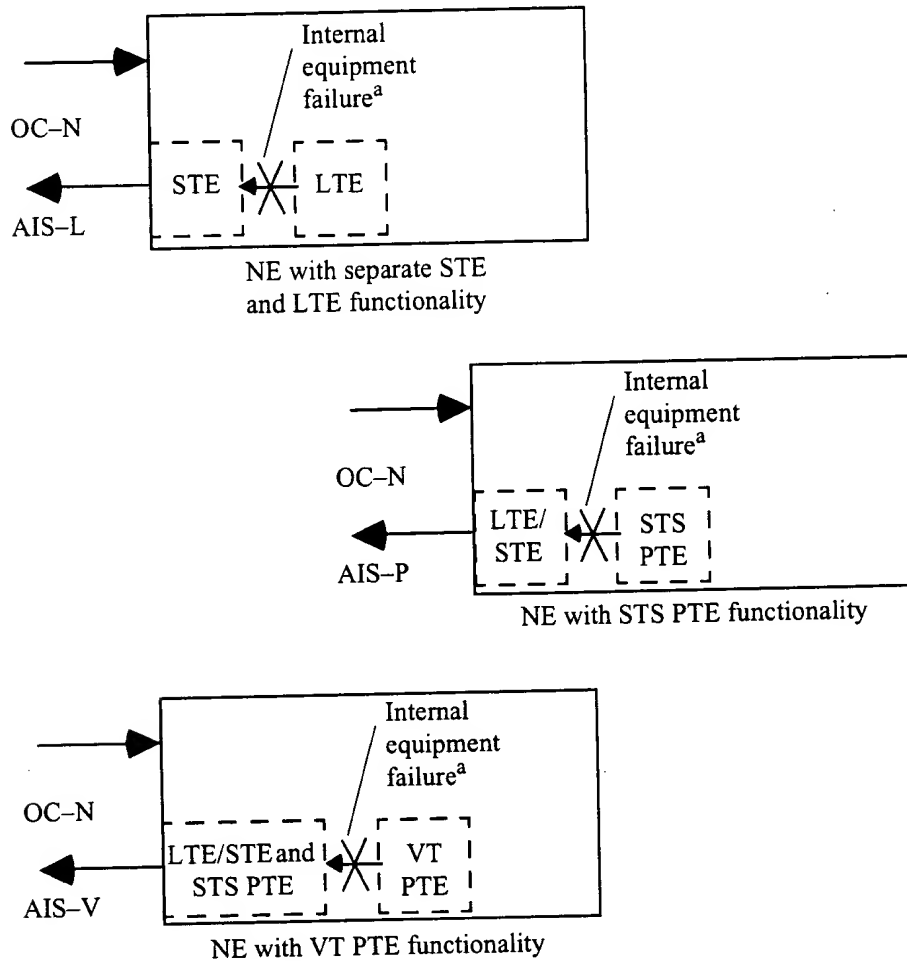


Figure 6-13. DS0 PTE (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning



Note:

a: In each case, the AIS is inserted when equipment that is originating a signal detects that the circuitry supporting provisioned higher layer origination functions has failed or been removed without having been "unprovisioned", and it continues until standby circuitry, if available, is switched in or the failure is cleared (see Section 6.2.1.2).

Figure 6-14. SONET Maintenance Signals for Internal Equipment Failures

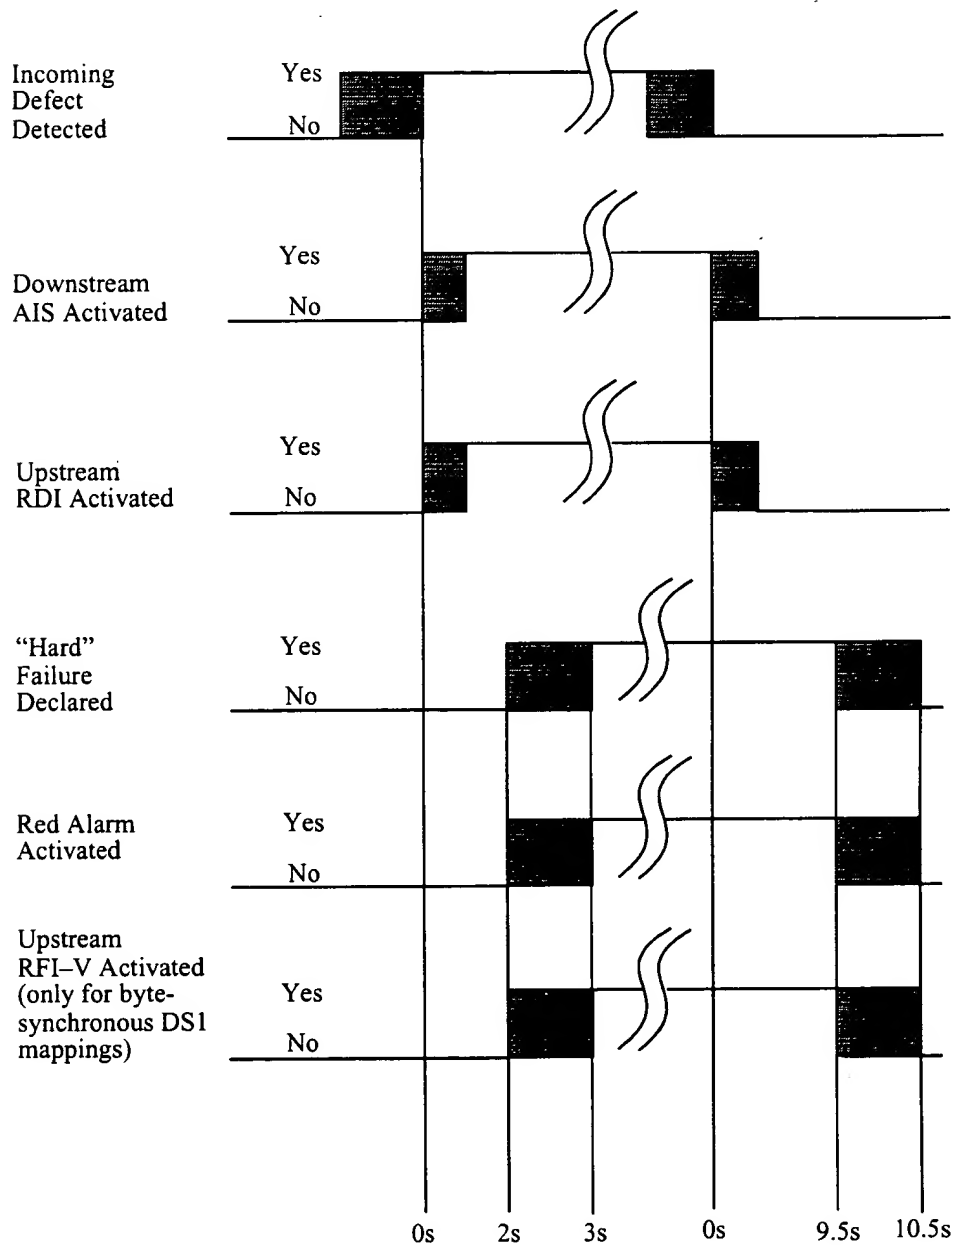


Figure 6-15. Alarm Timing Requirements for Directly Detected Defects and Failures

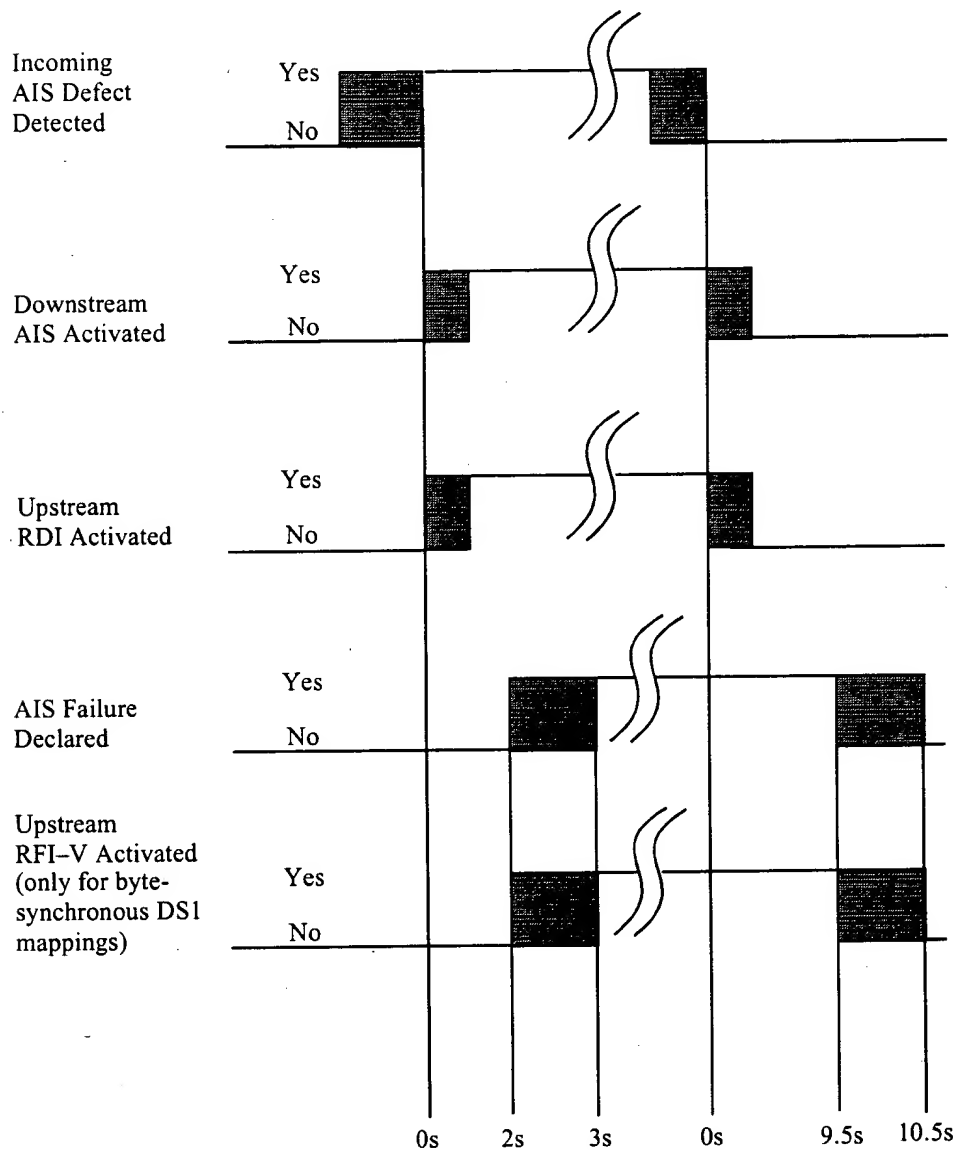


Figure 6-16. AIS Timing Requirements

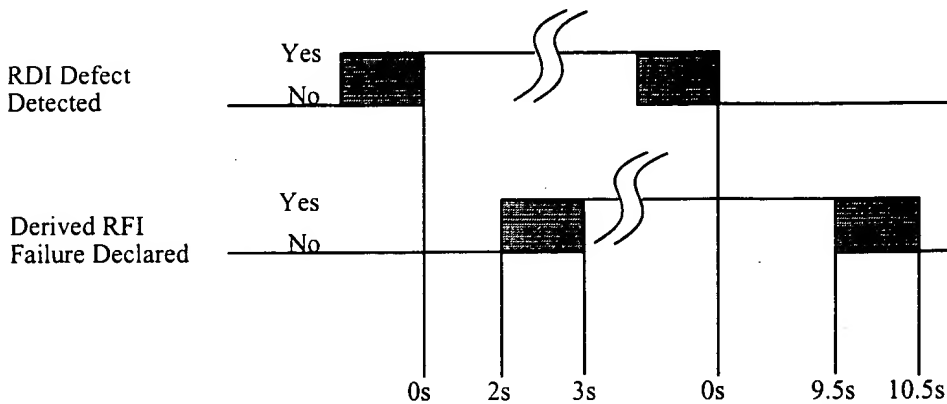


Figure 6-17. Derived RFI Timing Requirements

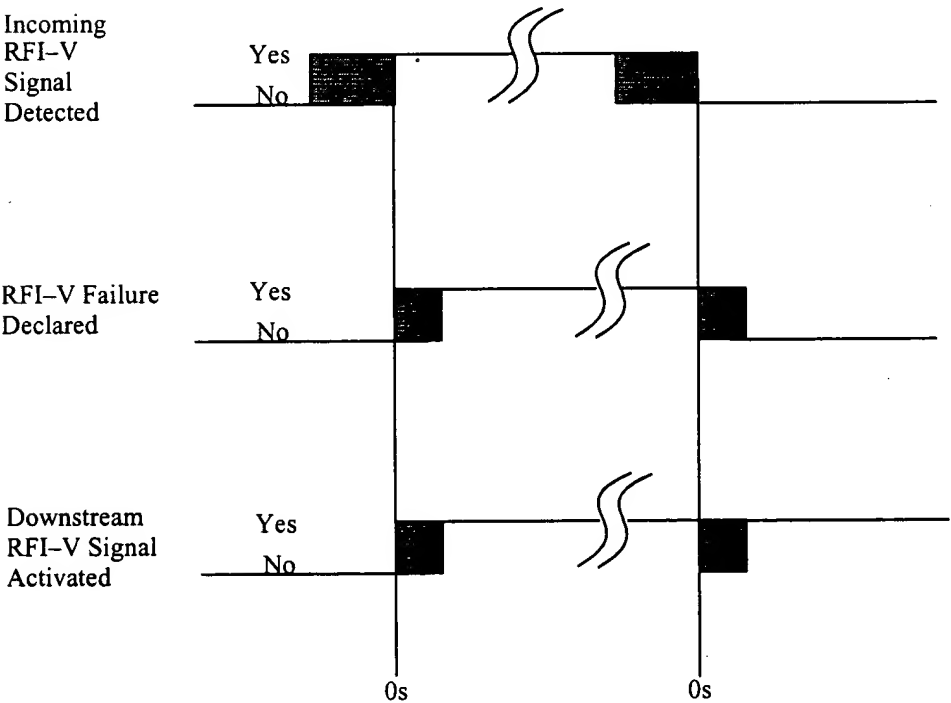


Figure 6-18. RFI-V Signal Timing Requirements (for Byte-Synchronous DS1 Mapping)

6.2.1.8 Control of Alarm Processing

GR-474-CORE contains criteria related to alarm processing that are applicable to transport and switching NEs in general. Although there are some differences in terminology and assumptions about the functionality of the NEs, many of the criteria in GR-474-CORE can be directly applied to SONET NEs. This section reviews and reiterates the criteria contained in GR-474-CORE that are applicable to transport NEs, and indicates where there are specific differences for SONET NEs.

In GR-474-CORE, the term "trouble notifications" is used to refer to the various ways that the user is alerted to the existence, location, and severity of a trouble. Trouble notifications include output messages (e.g., to an OS), visual indications at the NE, and audible/visible CO frame indications. The following sections are primarily concerned with the messages that an NE autonomously sends to an OS when it declares a failure, and the information that it sends in response to a user's request for the current condition of the NE.

Another document that contains information and criteria relevant to the area of alarm processing is GR-1093-CORE, *Generic State Requirements for Network Elements*, which contains information about the possible service states that may be supported by various NEs (including SONET NEs). In addition to the states currently discussed in that document, it has been proposed that an "Automatic In-Service Provisioning" (AISP) state also be defined. That state could be used during pre-service provisioning and testing of any type of termination point (e.g., DS1, OC-3). Based on the proposal, when a particular termination point is in the proposed state the NE would be required to perform most of the functions normally associated with that point (e.g., defect detection and termination, failure declaration and clearing, accumulation of PM data). However, it would not autonomously generate alarm and TCA messages for that termination point (for transmission to an OS). In addition, if a good signal were to be continuously received at the termination point for some provisionable time period, that point would be required to automatically transition to the In-Service state.

6.2.1.8.1 Alarm Level Designations

According to GR-474-CORE, a trouble notification message must contain certain information, including (among other items) the type of trouble, the time the event occurred (e.g., the time the failure was declared), the signal level affected, a designation as Service-Affecting (SA) or Non-Service-Affecting (NSA), a designation as alarmed or Not Alarmed (NA), and (for alarmed troubles) a designation as Critical (CR), Major (MJ), or Minor (MN). GR-474-CORE also provides defaults and guidelines regarding whether a trouble should be alarmed or NA, and whether alarmed troubles should be designated as CR, MJ or MN. In general, failures resulting from problems on incoming signals are supposed to be alarmed, while failures resulting from incoming maintenance signals are supposed to be NA or Not Reported. In addition, GR-474-CORE indicates that an SA failure that affects traffic equivalent to approximately five (or more) DS1s is supposed to

be designated as CR, an SA failure that affect one to five DS1s is supposed to be designated as MJ, and an alarmed NSA failure or SA failure affecting less than one DS1 is supposed to be designated as MN. In general, all of these criteria are applicable to SONET NEs; however, some additional criteria and explanations are needed.

For the user, the designation of a trouble as either SA or NSA is a very important issue. In general, an SA failure on an incoming signal is a non-AIS failure where the corresponding defect causes the detecting NE to generate AIS (or apply trunk conditioning) downstream, and the generated AIS appears at the NE's external interfaces after any available protection switching has been completed. For example, at an NE using 1+1 linear APS, an STS LOP failure declared on an STS-1 received on the selected line would be an SA failure, while an OC-N LOS would be considered NSA if the traffic is restored by a successful protection switch. Note that incoming AIS failures are generally not considered to be SA because the (upstream) NE that inserted the AIS should have already reported the root-cause incoming signal (or possibly equipment) problem as a SA failure. In addition, GR-474-CORE indicates that a declared AIS failure should either cause an NA notification or should be Not Reported, depending on the criteria in technology-specific document (such as this document). For AIS (and RFI) failures declared by SONET NEs, the following requirement applies:

R6-283 [621] AIS and RFI failures shall have a default setting of Not Reported.

R6-284 [622v2] The Far-End Protection Line failure shall have a default setting of a MN alarm.

6.2.1.8.2 *Single Failure/Single Message*

GR-474-CORE also indicates that any single failure (or root-cause incoming signal problem) must result in only one alarmed output message. This has a significant effect on SONET NEs, which can terminate multiple layers of the incoming signal, and for which multiple defects and failures are defined at each layer. The priorities (e.g., for alarm reporting purposes) of the various traffic-related failures that a SONET NE can declare based on its incoming SONET signals are discussed in this section.

In addition to traffic-related failures, a SONET NE may also monitor an incoming SONET signal for other types of failures. For example, a SONET signal could be carrying the APS channel, it could be used as a synchronization source, or it could be carrying operations communications via the Section DCC bytes. In general, the user may want to be alerted (via an alarmed output message) when one of those functions fails, even if that failure is a direct result of an incoming signal problem that causes a traffic-related alarmed output message (or an incoming maintenance signal).¹¹ Therefore, non-traffic-related failures may be exceptions to the "single failure/single message" criteria, and are not discussed in this section.

For the purposes of alarm reporting, traffic-related failures can be divided into near-end failures (i.e., LOS, LOF, LOP, AIS, UNEQ, TIM, and PLM failures), and far-end failures (i.e., RFI failures). In general, a near-end defect detected at a lower layer disrupts the capability of the equipment terminating higher layers to monitor for both near-end and far-end defects, unless a successful protection switch restores the signal to the higher layer. For example, an LOF defect detected by an NE's STE disrupts the LTE's ability to monitor for an externally generated AIS-L, RDI-L and (if the LTE processes STS pointers) LOP-P defects, but does not disrupt the STS PTE's ability to monitor for STS path layer defects if a line-level protection switch is successfully completed. Therefore, the failure resulting from the lowest-layer defect detected is the one reported to the OS.

For near-end defects that are detected at the same layer, one of the following applies:

- The terminating equipment monitors the same part of the incoming signal to detect and terminate two different defects, and therefore the defects are defined such that the detection of one of them is a condition for terminating the other one. In some cases the defects have been defined to result in symmetric operation (e.g., the detection of an UNEQ-P defect terminates a PLM-P defect, and the detection of any "non-unequipped" signal label code, including one that results in a PLM-P defect, terminates an UNEQ-P defect), while others may be asymmetric (e.g., the detection of an AIS-P defect terminates an LOP-P defect, but if an NE does not meet **O6-174 [1079]** then the AIS-P defect will only be terminated if a "valid" pointer is detected).
- One of the incoming defects disrupts the ability of the terminating equipment to access its POH (i.e., one of the incoming defects is an LOP or AIS defect). In such cases, the defects that the PTE would normally detect or terminate based on the contents of the POH [i.e., UNEQ, TIM, PLM (and RDI)] cannot be detected or terminated.

Based on the preceding discussion, an NE that uses internal SONET signals between layers, that has an alarm level of Not Reported for AIS failures, and that meets the existing criteria on detecting defects and declaring failures would be expected to automatically generate a single alarmed (traffic-related) output message for almost any single incoming signal problem. However, for other cases (e.g., for a TIM-P or PLM-P failure that is declared in response to the same incoming signal problem as an UNEQ-P or TIM-P failure, or where AIS failures are set to be alarmed) the following requirement is necessary.

R6-285 [626v2] A SONET NE shall not autonomously report a near-end failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE, per the hierarchy in Table 6-6. In addition, the SONET NE shall not autonomously report a near-end failure declared for equipment (e.g., STS PTE) that has been

11. Note however, that to avoid the declaration of unnecessary alarms when the LTE terminating the protection line in a linear APS system detects an AIS-L (or lower layer) defect, the APS-related defects are defined such that they will not be detected when an AIS-L defect is present.

provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8).

Table 6-6. Hierarchy of Near-end Failures

Priority	Failure
Highest	LOS
	LOF
	AIS-L
	AIS-P ^a
	LOP-P ^b
	UNEQ-P
	TIM-P
	PLM-P
	AIS-V ^a
	LOP-V ^b
	UNEQ-V
	PLM-V
Lowest	DSn AIS (if Reported for outgoing DSn signals)

Notes:

- Although it is not defined as a defect/failure, all-ones STS pointer relay is also higher priority than LOP-P. Similarly, all-ones VT pointer relay is higher priority than LOP-V.
- LOP-P is also higher priority than the far-end failure RFI-P, which does not affect the detection of any of the near-end failures. Similarly, LOP-V is higher priority than RFI-V.

For far-end defects (i.e., RDI defects), the detection of a defect at a particular layer does not disrupt the capability to detect defects (either near-end or far-end) at any higher layer. Therefore, it is possible for an NE to detect multiple RDI defects and declare multiple RFI failures simultaneously. In addition, the detection of a single near-end defect by the far-end NE could cause that NE to generate multiple upstream RDI signals. If the near-end NE is set to report RFI failures, then that single root-cause incoming signal problem or maintenance signal (at the far-end) could result in multiple failure messages being sent to the OS by the near-end NE. To avoid this, the required default setting for RFI failures is Not Reported. In addition, for NEs that are set to report RFI failures, the following requirement applies.

- R6-286 [629v3]** An NE that is set to report RFI failures shall not autonomously report an RFI failure that is apparently caused by the same root-cause incoming signal problem or maintenance signal at the far-end that caused the NE to concurrently declare (and report) a higher-priority RFI failure, per the hierarchy in Table 6-7. In addition, the SONET NE shall not

autonomously report a far-end failure declared for equipment (e.g., STS PTE) that has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8).

Table 6-7. Hierarchy of Far-end Failures

Priority	Failure ^a
Highest	RFI-L
	RFI-P
Lowest	RFI-V

Notes:

- a. Note that Issue 1 of this document also included the Far-end Protection Line failure in the hierarchy of far-end failures, between RFI-L and RFI-P. That failure has been removed because it is considered an independent, APS-related failure; however, an NE may include it in the hierarchy and still conform to the requirement.

For example, an NE that is using linear APS and that detects RDI-L defects on both the working and the protection lines from an adjacent NE, along with RDI-P and RDI-V defects on all of the terminating STS and VT paths received on those lines would only report the RFI-L failures. It would not report all of the lower priority RFI-P and RFI-V failures.

6.2.1.8.3 Independent Failures

Section 6.2.1.8.2 is concerned with the failures reported to an OS due to a single root-cause incoming signal problem or maintenance signal, while this section is concerned with independent failures, as defined below.

- Independent failures - Failures (as defined in Section 6.2.1) that are declared for a single entity (i.e., a SONET Section, Line, STS Path, VT Path, or possibly an outgoing DS_n) in response to two or more root-cause incoming signal problems or maintenance signals, possibly requiring separate maintenance actions.

In general, it is assumed that for an NE to determine that two failures are independent, those failures would need to be declared at "separate times". However no assumptions have been made regarding a definition of "separate times", as this is viewed as an implementation detail. (One possible definition is that two failures are considered independent if the first failure has already been declared at the point in time when the defect leading to the second failure is detected.)

As an example of the above definition, if a fiber break in the optical path carrying an unprotected OC-3 signal causes an NE to (approximately simultaneously) declare an OC-3 LOS failure, 3 AIS-P failures and 84 AIS-V failures, those failures would not be

considered independent (and the criteria in Section 6.2.1.8.2 would apply). However, if prior to the fiber break, the NE had declared an UNEQ-V failure for one of the VT paths, then that UNEQ-V failure and the subsequent AIS-V failure on that same path would be considered to be independent (and the criteria in this section would apply).

Issue 1 of this document contained several requirements that indicated that certain failures were supposed to be cleared when certain other failures were declared. It also contained requirements indicating that certain defects were supposed to be terminated when certain higher-priority defects were detected, even though the higher-priority defects disrupted the signal such that the bits or bytes that the NE needed to monitor to terminate the lower-priority defect could not be accessed. Those requirements covered only a few of the possible combinations of defects and failures, and could cause confusion for combinations that were not covered. In addition, failures that are declared at separate times are generally caused by separate problems, which need to be individually addressed. Therefore, those requirements have been removed, and **R6-287 [988]** (which is currently under review in GR-253-ILR, Issue ID 253-121) has been added. While the former requirements are no longer considered necessary, an NE that meets those former requirements should not be considered nonconforming to the current criteria. Therefore, the following are allowed exceptions to **R6-287 [988]** and the defect termination criteria in Sections 6.2.1.1.8 and 6.2.1.3:

- An NE may clear an existing LOF failure when an LOS failure is declared
- An NE may terminate an UNEQ-P defect when an AIS-P or LOP-P defect is detected
- An NE may clear an UNEQ-P failure when an AIS-P or LOP-P failure is declared
- An NE may terminate a PLM-P defect when an AIS-P or LOP-P defect is detected
- An NE may clear a PLM-P failure when an AIS-P or LOP-P failure is declared
- An NE may terminate an RDI-P defect when an AIS-P defect is detected
- An NE may terminate an UNEQ-V defect when an AIS-V or LOP-V defect is detected
- An NE may clear an UNEQ-V failure when an AIS-V or LOP-V failure is declared
- An NE may terminate a PLM-V defect when an AIS-V or LOP-V defect is detected
- An NE may clear a PLM-V failure when an AIS-V or LOP-V failure is declared
- An NE may terminate an RDI-V defect when an AIS-V defect is detected

R6-287 [988] The declaration of a "new" failure by a SONET NE shall not automatically cause the NE to clear any previously declared, independent failures.

In some cases, an NE that meets the preceding requirement would still be expected to clear an existing failure after a new higher priority (or lower layer) failure is declared; however,

the existing criteria on defect detection and termination, and failure declaration and clearing should be sufficient. In other cases, the lower priority failure would be expected to remain in place because the defect that caused it cannot be terminated while the defect associated with the higher priority failure is present.

6.2.1.8.4 *Retrieval of NE Condition*

While Sections 6.2.1.8.2 and 6.2.1.8.3 were concerned with the autonomous messages that an NE sends to an OS, this section discusses the response of the NE to requests from the user for a report on the current condition of the NE.

In general, when an NE is responding to a user request to report all of the failures at the NE, it needs to report all of the independent failures (i.e., the same failures that it has autonomously reported). If the NE were to also report lower-priority failures that were caused by the same root-cause incoming signal problem or maintenance signal as a higher-priority failure, then the number of failures reported to the user could easily become excessive. On the other hand, if the user requests only the failures at a particular layer or for a particular entity (e.g., for the PTE supporting a particular path), then it would be appropriate for the NE to report all of the failures that it has declared for that layer or entity.

R6-288 [625] A SONET NE shall provide the capability to report, on-demand to the user, the current failure indications (i.e., the current condition of the NE).

R6-289 [627v2] A SONET NE shall not report a failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE (per the hierarchy in Table 6-6) in response to a request to report all failures at the NE.

R6-290 [628v2] A SONET NE shall report all failures, including each failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE, in response to a request to report all failures at a given SONET layer, or for a given entity.

For example, if an AIS-L failure is declared concurrently with an AIS-P failure, indications for both failures would be kept at the NE, even though only the AIS-L failure would be autonomously reported. In that case, if the NE were subsequently requested by the user to report all failures, it would report the AIS-L failure (along with any other independent failures). However, if a request was received to report STS path layer failures, it would report the AIS-P failure (along with any other STS path layer failures).

6.2.1.8.5 *Provisioning of Alarm Levels*

GR-474-CORE also contains an objective for the designation of a trouble as alarmed or NA to be provisionable, and a requirement that the user be able to inhibit any NA notification (e.g., by changing the designating to Not Reported). In SONET, the capability to provision certain failures as alarmed or NA is changed from an objective to the following requirement.

R6-291 [620] SONET equipment shall provide the capability of setting any AIS (including DS_n AIS), RFI (including DS_n RAI), and Far-End Protection Line failure as either Reported or Not Reported, and if Reported, as either alarmed or Not Alarmed. The settings shall be provisionable on a per-layer, per-entity basis (e.g., for the line layer, the settings shall be provisionable on a per-line basis).

R6-292 [623] A SONET NE shall provide the capability of reporting (on demand) all software settable attributes.

GR-474-CORE contains requirements for controlling local alarms (e.g., alarm cutoff) and remote alarms.

6.2.1.8.6 *Clear Messages*

Finally, GR-474-CORE indicates that a clear message must be generated when an alarmed trouble clears. However, it does not indicate that a clear message also needs to be generated when certain NA troubles that were reported to the OS are cleared [i.e., when NA troubles that were reported as standing conditions (as opposed to transient conditions) are cleared]. Therefore, the requirement on clearing failures is expanded to the following for SONET NEs:

R6-293 [632] A SONET NE shall individually clear (and send a clear message to the OS) any failure that is individually reported to an OS.

6.2.1.8.7 *Non-Intrusive Detection of Defects and Declaration of Failures*

In some applications, equipment within an NE may need to detect defects that, based on the criteria in Sections 6.2.1.1.1 through 6.2.1.6, are normally detected only by equipment with other (higher-level) functionality. For example, STS PTE that generates PDI-P for an outgoing VT-structured STS-1 must detect AIS-V defects even though it does not terminate the VT path. Based on wording of the AIS-V detection criteria in Section 6.2.1.2.3, those criteria are applicable only to VT PTE. Therefore, the PDI-P criteria in Section 6.2.1.4.1 indicate that the STS PTE "shall detect and terminate AIS-V

defects as if it were VT PTE". Similar statements are made in other criteria, and in general, the equipment that detects and terminates the defects does so non-intrusively. The equipment may take some action based on the detection of the defect (e.g., change the PDI-P code that it is sending), but it does not alter the monitored signal by generating (rather than passing AIS) downstream.

When an NE non-intrusively monitors a signal, it is generally not appropriate for it to autonomously report (to an OS) any failures that it could declare based on the defects detected on that signal.¹² The reason for this is that when the signal is terminated at a downstream NE, that NE will generally detect the same defect, declare the same failure, and report that failure to an OS (unless it is a Not Reported failure). If a number of NEs non-intrusively monitor a signal and report failures on that signal to an OS, the OS could receive a large number of reports for a single failure.

Although in most applications it is not appropriate for an NE that non-intrusively monitors a signal to autonomously report failures related to that signal to an OS, some users may want that capability. In addition, it can be useful for the NE to declare and clear those failures so that they can be retrieved by the user (e.g., for trouble-shooting purposes). Therefore, the following criteria are applicable:

R6-294 [695v2] An NE that is non-intrusively monitoring a signal shall declare and clear failures for that signal as if it were terminating the signal. Upon declaring or clearing a failure, the NE shall set or clear the appropriate failure indication for that signal.

O6-295 [703v3] As a default, an NE that is non-intrusively monitoring a signal should not autonomously report to an OS the declaration or clearing of a failure on that signal (i.e., the default or "fixed" setting for the failures should be Not Reported).

Note that although **O6-295** [703v3] refers to a "default" setting for failures declared for non-intrusively monitored signals, this is not meant to imply that the levels for those failures are required to be provisionable. Instead, it is intended to indicate that an NE is allowed to support provisionable alarm levels for those failures, and that if an NE does happen to support them, then the default level should be Not Reported. Similarly, if the NE does not happen to support provisionable alarm levels for those failures, then the "fixed" (or only supported) level should be Not Reported.

12. One exception to this may be defects that an NE uses for path protection switching purposes. See the appropriate application-specific GRs (e.g., GR-1230-CORE, GR-1400-CORE) for any additional criteria related to the reports generated when an NE detects such a defect and performs a path protection switch.

6.2.2 Performance Monitoring

Performance Monitoring (PM) refers to the in-service, non-intrusive monitoring of transmission quality. A SONET NE is required to support PM according to the layer of functionality that it provides, and in SONET the capability is provided to accumulate PM data based on overhead bits (e.g., BIP-Ns) at the Section, Line, and STS Path, and VT Path layers. In addition, PM data is available at the SONET Physical layer using physical parameters (rather than overhead bits). For SONET NEs that interface to the existing digital network (e.g., DS1), the accumulation of certain DS_n PM parameters may also be appropriate.

R6-296 [633] Except as specifically noted, SONET NEs shall meet the general PM requirements in GR-820-CORE.

GR-820-CORE contains generic PM strategies, discusses various types of PM registers (e.g., current period, previous period, recent period, and threshold registers), and defines PM parameters for DS1 and DS3 signals. In general, the generic strategies discussed in GR-820-CORE are directly applicable for monitoring SONET signals, and therefore are not repeated here. However, additional PM parameters are needed to support SONET PM. The following sections describe the PM parameters defined for each SONET layer, and identify the accumulation and thresholding requirements for SONET NEs. At certain SONET layers, both near-end and far-end performance can be monitored. For those layers, both near-end and far-end PM parameter definitions, accumulation requirements, and thresholding requirements are provided.

In general, an NE accumulates various PM parameters based on performance “primitives” that it detects in the incoming digital bit stream. Primitives can be either “anomalies” or defects. An anomaly is defined to be a discrepancy between the actual and desired characteristics of an item (e.g., an error in a received BIP-8 code), and may or may not affect the ability of the item to perform a function. A defect is defined to be a limited interruption in the ability of an item to perform a required function. The persistence of a defect results in a failure, which is defined to be the termination of the ability of an item to perform a required function (see Sections 5.5 and 6.2.1 for descriptions of the various defects and failures). Functionally, PM is performed at each layer, independent of the other layers. However as Figure 6-2 illustrates, part of the functional model assumes that layers pass maintenance signals to higher layers. For example, a defect (e.g., LOS) that occurs at the Section layer causes an AIS-L signal to (functionally) be passed up to the Line layer, which causes the AIS-P signal to be passed up to the STS Path layer, etc. Thus, an AIS defect can be detected at a particular layer either by receiving the appropriate AIS on the incoming signal, or by receiving it from a lower layer. Therefore, defects and failures that occur at a lower layer affect the PM parameters at higher layers, and the definitions of the assorted PM parameters were written with that consideration.

Thresholds are defined for most of the PM parameters supported by SONET NEs, and are used to detect when transmission degradations have reached unacceptable levels. When a

threshold for a non-Physical layer parameter is reached or exceeded and there is no potential for the parameter to be adjusted to be less than the threshold because of entry into (or exit from) unavailable time,¹³ a Threshold Crossing Alert (TCA) is sent to an OS. Similarly, when the recorded value of a Physical layer parameter is outside of its acceptable range, an out-of-range alert message is sent to an OS. A TCA or out-of-range alert is reported as a transient condition (i.e., it is not cleared at some later time), and only one TCA or out-of-range alert can be generated based on the data accumulated or recorded in any particular current period register. However, it is possible for another TCA or out-of-range alert to be generated based on the data accumulated or recorded in the next accumulation period (e.g., if the threshold is again reached or exceeded).

Figures 6-19, 6-20, and 6-21 illustrate (in the form of flowcharts) some of the processes involved in the accumulation of non-Physical layer PM data. Each process runs concurrently with the others, although some are triggered by the same conditions. These figures are intended to clarify the requirements in the following sections, and therefore they should be viewed only as a logical representation of actions that could be taken to meet the requirements. They are not intended to imply or unnecessarily constrain an implementation. In fact, some steps are intentionally vague or omitted, because they are considered implementation details.

In addition to the types of registers discussed in GR-820-CORE, several other types of registers are mentioned in Figures 6-19, 6-20, and 6-21. Note that in all cases, "register" simply means a location where data is stored, and should not be taken to imply any type of architectural design. In the flowcharts, the term "current second register" refers to a register that contains only the data for the current second. At the end of a second, the data in the current second register is normally moved to the current period registers, unless some other action is warranted. Each parameter at the Section, Line, and Path layers has both a current second register and current period registers. Also, "negative adjustment registers" and "positive adjustment registers" (referred to jointly as "inhibition registers" in Issue 1 of this document) refer to registers that are used to subtract off or add to a current period register (and in some cases, possibly the previous period register) when unavailable time is entered or exited. Each Coding Violation (CV), Errored Second (ES), Severely Errored Second (SES), Unavailable Second (UAS), Pointer Justification (PJ) parameter at the Line, STS Path, and VT Path layers has a negative adjustment register, a positive adjustment register, or both.

The following list describes the purpose of each of the flowcharts (which are labeled A through J):

- A: Illustrates how a current second register is incremented
- B: Illustrates how a current period register is accumulated

13. See Sections 6.2.2.4 through 6.2.2.6 for the definitions of unavailable time at the SONET Line, STS Path, and VT Path layers, and Section 6.2.2.8 for criteria on the accumulation of PM parameters during unavailable time (including entry into and exit from unavailable time).

C: Illustrates conditions under which a TCA message is sent

D: Illustrates how PM registers behave as a push-down stack

E and F: Illustrate cases where the negative and positive adjustment registers are incremented

G and H: Illustrate cases when the negative and positive adjustment registers are cleared

I: Illustrates how current period registers are adjusted on entering unavailable time

J: Illustrates how current period registers are adjusted on exiting unavailable time

Note that not all of the flowcharts are applicable for all parameters, and that the flowcharts shown do not form a complete set. For example, flowchart C does not cover several cases where a TCA may need to be sent during the first ten seconds of period based on the data accumulated for the previous period (see Section 6.2.2.1). In addition, flowcharts C and E through J are not applicable for UAS parameters (for which TCAs would normally be sent during unavailable time rather than available time, and the roles of the negative and positive adjustment registers are reversed), flowcharts E through J are not applicable for the parameters at the Section layer (at which no UAS parameter has been defined), and flowcharts E, G and I are not applicable for Line or Path CV parameters (the accumulation of which is already inhibited during the first ten seconds of unavailable time because those seconds are SESs).

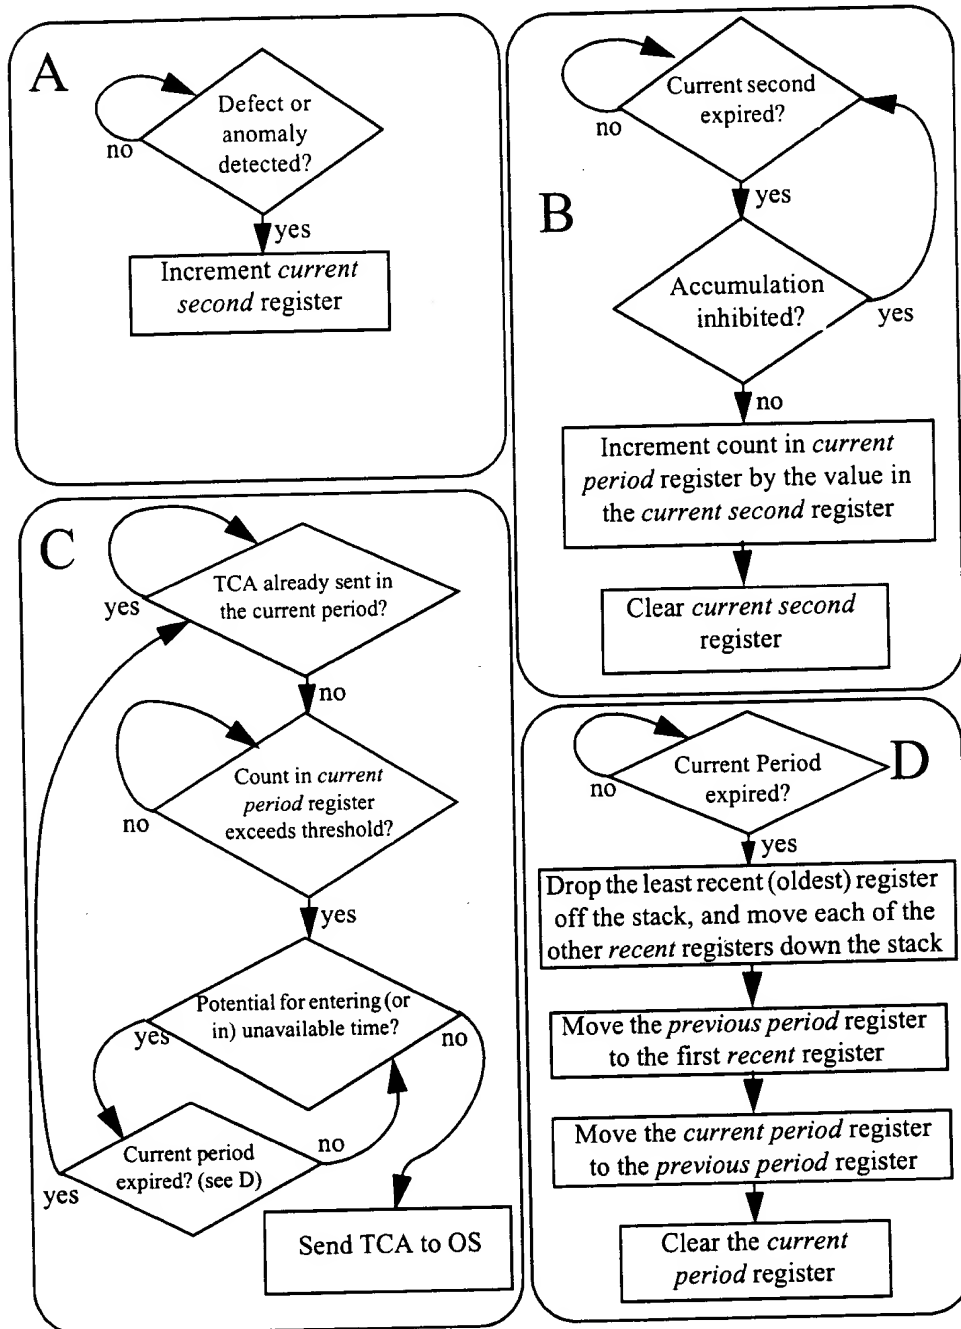


Figure 6-19. SONET PM Accumulation and Thresholding Model

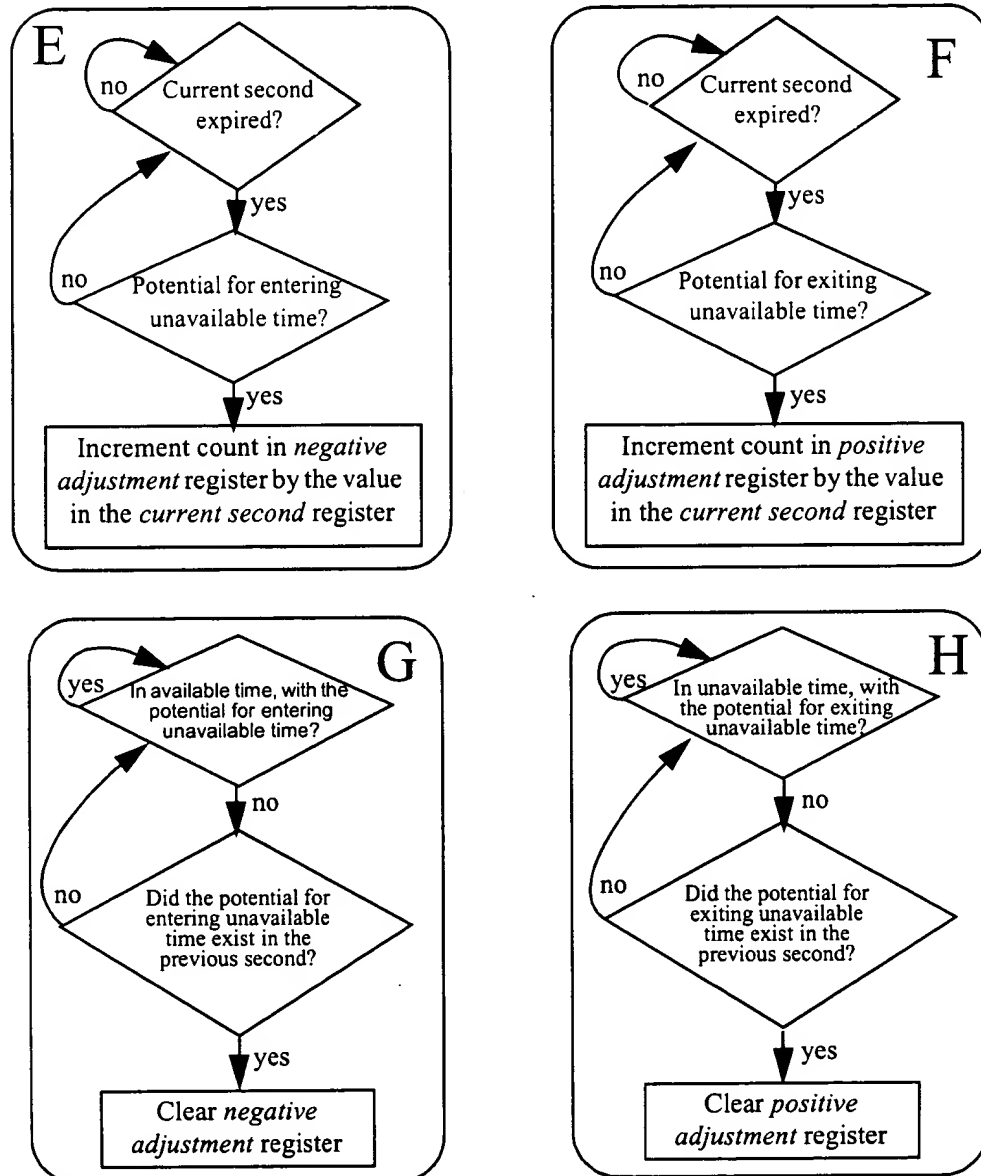


Figure 6-20. SONET PM Accumulation and Thresholding Model (Continued)

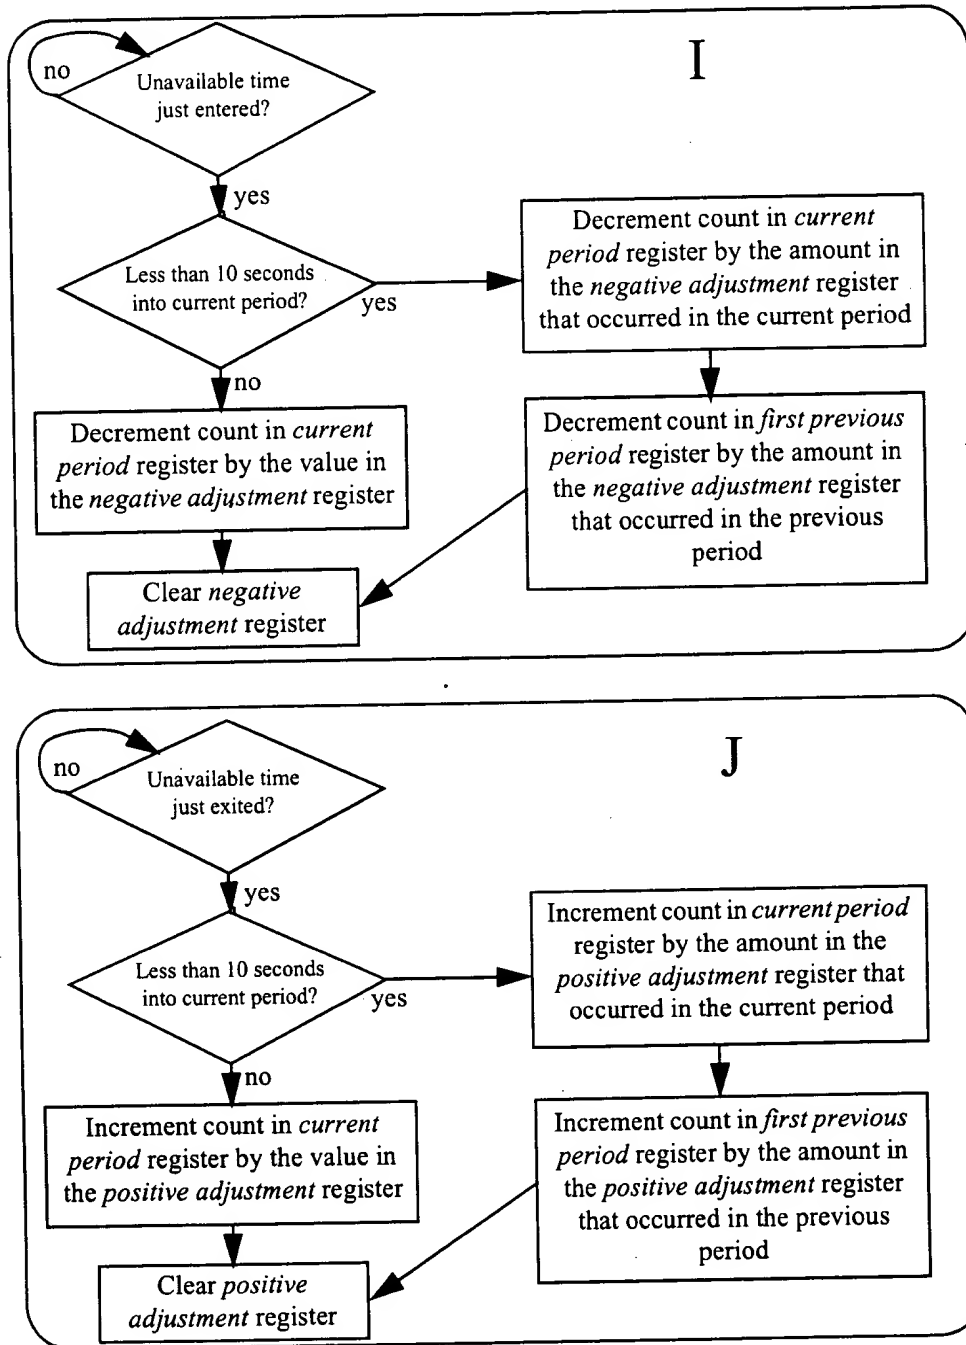


Figure 6-21. SONET PM Accumulation and Thresholding Model (Continued)

6.2.2.1 General Accumulation and Thresholding Criteria

SONET NEs are required to accumulate a variety of PM parameters. Requirements indicating the number and types of registers needed for each parameter are presented below. The wording of the criteria are meant to illustrate the requirements, not to carry any design implication.

- R6-297** [634v2] For each PM parameter accumulated for a Physical, Section, Line, STS Path or VT Path layer entity, a SONET NE shall provide one current 15-minute, one current day, one previous 15-minute, one previous day, and 31 recent 15-minute accumulation and storage registers.
- R6-298** [639v2] The size of the PM parameter accumulation registers provided by a SONET NE shall be greater than or equal to the minimum accumulation register sizes shown in Table 6-8.
- R6-299** [636] A SONET NE shall allow the user to initialize all current 15-minute or current-day registers to zero at any time, on an individual entity (e.g., STS Path) basis, per direction (e.g., near-end).
- R6-300** [637v2] At the end of each 15-minute period, the current 15-minute register, the previous 15-minute register, and the 31 recent 15-minute registers shall behave as a push-down stack. The current 15-minute registers shall then automatically be initialized to zero.

In a push-down stack, the contents of the least recent 15-minute register is pushed off the stack, each of the thirty remaining recent 15-minute registers is pushed down the stack, the previous 15-minute register is moved to the most recent 15-minute register, and the current 15-minute register is moved to the previous 15-minute register. The (new) current 15-minute register is then initialized to zero.

- R6-301** [989] At the end of each day, the contents of each current-day register shall be copied to the corresponding previous day register. The current day registers shall then automatically be initialized to zero.

The invalid-data flag is set to indicate that the data stored in a register or set of registers may be corrupt, because either the period of accumulation is greater than or less than the nominal period (e.g., the NE's time of day setting was changed during the period), or data is missing during an accumulation period (e.g., one or more of the current registers were initialized during the period, the NE was unable to perform its PM data accumulation functions for some or all of the period, or the accumulation of the far-end parameters was inhibited due to the presence of a near-end defect for some or all of the period).

- R6-302** [635v3] A SONET NE shall provide, as a minimum, an invalid-data flag associated with each monitored entity's near-end parameters, and another invalid-data flag associated with its far-end parameters (if applicable).

Note that it is also acceptable for an NE to provide separate invalid-data flags for each register.

- R6-303** [990] Invalid-data flags shall be moved with the data to which they apply.

In addition to the accumulation and storage registers required for most PM parameters, a SONET NE must also provide threshold registers according to the following requirements.

- R6-304** [638v3] The following apply regarding threshold registers:

- ... • 15-minute and 1-day threshold registers shall be provided for the SONET PM parameters that require thresholding (see Sections 6.2.2.2 through 6.2.2.7)
- ... • The size of the threshold registers shall be greater than or equal to the minimum threshold register sizes shown in Table 6-8
- ... • The value in each threshold register (with the possible exception of certain Physical layer threshold registers, see Section 6.2.2.2.2) shall be provisionable
- ... • Default values for all threshold registers shall be provided and documented by the NE supplier
- ... • Where equivalent near-end and far-end PM parameters are defined, the default threshold value shall be the same for the near-end and the far-end parameters
- ... • For PM parameters where default threshold values are shown in Table 6-8, the default values provided by the NE shall be equal to those values.

Note that where equivalent near-end and far-end PM parameters are defined, it is sufficient to provide a single threshold register for both. An NE may optionally provide separate threshold registers for the near-end and far-end parameters, along with the capability to provision different threshold values for those parameters.

- R6-305** [1081] Unless the applicable equipment (e.g., the optical transmitter) has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8.2), an out-of-range alert shall be sent to an OS when the new "snapshot" value in a current Physical layer 15-minute or current day register is not in the acceptable range for that parameter [i.e., when it is less than or equal to the lower threshold (if defined), or is greater than or equal to the upper threshold].

R6-306 [991v2] Unless the applicable equipment (e.g., the LTE) has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8.2), a TCA shall be sent to an OS when the value in a current non-Physical layer 15-minute or current day register reaches or exceeds the corresponding threshold and it is not possible for that value to be adjusted to less than the threshold based on entry into or exit from unavailable time.

Based on the definitions of the UAS parameters at the Line and Path layers and the requirements on the accumulation of various parameters during unavailable time, an NE must be capable of adjusting the data in certain previous period registers during the first ten seconds of a new period (due to entry into or exit from unavailable time). In addition, such an adjustment could cause the value in the previous period register to reach or exceed the corresponding threshold, in which case the NE would need to send a TCA for the previous period. Similarly, an NE would need to send a TCA for the previous period if the value in a register has reached or exceeded its threshold, a TCA has not been sent because of the potential for entry into or exit from unavailable time, the period expires (so the data in the current register is moved to the previous period register), and then unavailable time is not entered or exited.

While the data adjustment and TCA generation capabilities described above are already required to be provided based on other criteria, it is necessary to strictly interpret and correlate those criteria in order to determine that they are required. Therefore, the following requirement explicitly addresses those capabilities. Note that this is applicable only for parameters whose accumulation is affected in some way by entry into or exit from unavailable time.

R6-307 [992v2] The following apply regarding the adjustment of PM data in previous period registers (due to entry into or exit from unavailable time) and the generation of TCAs.

- ...
 - An NE shall either provide the capability to adjust the data stored in its most recent previous period registers based on entry into or exit from unavailable time during the first 10 seconds of a new current period, or it shall delay updating the values in its current period registers (and the generation of TCAs and the moving of data between registers) for up to 10 seconds.
- ...
 - If the capability to adjust the data stored in previous period registers is supported, a TCA shall be sent to an OS when the value in a previous 15-minute or previous day register is adjusted so that it is greater than or equal to the corresponding threshold.
- ...
 - If the capability to adjust the data stored in previous period registers is supported, a TCA shall be sent to an OS when the value in a previous 15-minute or previous day register can no longer be adjusted so that it

will be less than the corresponding threshold (i.e., when the potential entry into or exit from unavailable time that was inhibiting the generation of the TCA at the end of a period does not occur).

- If the NE delays updating its registers (and the generation of TCAs and the moving of data between registers) for up to 10 seconds, it shall use that delay time to determine if anything has occurred that should affect the data about to be reflected in the current period registers.

In general, an NE that uses a short delay as described above would still be considered to be conforming to the objective in GR-820-CORE to accumulate PM data in "real-time".

In addition to the autonomous transmission of TCA (and out-of-range alert messages) by an NE, a user may also query the NE for the values in any of its PM registers, including its current period registers and threshold registers.

- R6-308** [640v3] The SONET NE shall provide the capability for the user to retrieve the contents of any PM parameter register at any time. It shall also provide the capability to retrieve the contents of any threshold register (i.e., the provisioned threshold value).
- R6-309** [641] The SONET NE shall allow the user to schedule, and shall then perform periodic (and automatic) reporting of PM data for a monitored entity. The NE shall continue to send the appropriate PM data according to the schedule until instructed to stop by the user. This instruction to stop could be part of the scheduling information that started the periodic reporting, or it could be a separate request. The NE shall support both methods of stopping periodic performance reporting.
- R6-310** [642] The SONET NE shall support the ability for the user to retrieve periodic PM report schedule information.

Table 6-8. PM Register Sizes and Default Thresholds

Parameter	Rate (e.g., OC-N, STS-M)	Minimum Register Size		Minimum Threshold Register Size		Threshold Default Value	
		15 minute	1 day	15 minute	1 day	15 minute	1 day
Physical							
LBC _{normal}	Optical	-	-	255 ^a	255 ^a	-	-
OPT _{normal}	Optical	-	-	255 ^a	255 ^a	-	-
OPR _{normal}	Optical	-	-	255 ^a	255 ^a	-	-
Section							
CV-S	N = 1, 3	16,383	1,048,575	16,383	1,048,575	-	-
	12	16,383	1,048,575	16,383	1,048,575	-	-
	24	16,383	1,048,575	16,383	1,048,575	-	-
	48	16,383	1,048,575	16,383	1,048,575	-	-
ES-S	All	900	65,535	900	65,535	-	-
SES-S	All	900	65,535	900	65,535	-	-
SEFS-S	All	900	65,535	900	65,535	-	-
Line							
CV-L CV-LFE	N = 1, 3	16,383	1,048,575	16,383	1,048,575	-	-
	12	16,383	1,048,575	16,383	1,048,575	-	-
	24	16,383	1,048,575	16,383	1,048,575	-	-
	48	16,383	1,048,575	16,383	1,048,575	-	-
ES-L ES-LFE	All	900	65,535	900	65,535	-	-
SES-L SES-LFE	All	900	65,535	900	65,535	-	-
UAS-L UAS-LFE	All	900	65,535	900	65,535	-	-
FC-L FC-LFE	All	72	4,095	NR	NR	NR	NR
PSC	All	63	255	NR	NR	NR	NR
PSD	All	900	65,535	NR	NR	NR	NR

Table 6-8. PM Register Sizes and Default Thresholds (Continued)

Parameter	Rate (e.g., OC-N, STS-M)	Minimum Register Size		Minimum Threshold Register Size		Threshold Default Value	
		15 minute	1 day	15 minute	1 day	15 minute	1 day
STS Path							
CV-P	M = 1	16,383	1,048,575	16,383	1,048,575	15	125
CV-PFE	3c	16,383	1,048,575	16,383	1,048,575	25	250
	12c	16,383	1,048,575	16,383	1,048,575	75	750
ES-P	M = 1	900	65,535	900	65,535	12	100
ES-PFE	3c	900	65,535	900	65,535	20	200
	12c	900	65,535	900	65,535	60	600
SES-P	All	900	65,535	900	65,535	3	7
SES-PFE							
UAS-P	All	900	65,535	900	65,535	10	10
UAS-PFE							
FC-P	All	72	4,095	NR	NR	NR	NR
FC-PFE							
PPJC-PDet	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
NPJC-PDet							
PPJC-PGen	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
NPJC-PGen							
PJCDiff-P	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
PJCS-PDet	All	900	65,535	900	65,535	-	-
PJCS-PGen							

Table 6-8. PM Register Sizes and Default Thresholds (Continued)

Parameter	Rate (e.g., OC-N, STS-M)	Minimum Register Size		Minimum Threshold Register Size		Threshold Default Value	
		15 minute	1 day	15 minute	1 day	15 minute	1 day
VT Path							
CV-V CV-VFE	All	16,383	1,048,575	16,383	1,048,575	-	-
ES-V ES-VFE	All	900	65,535	900	65,535	-	-
SES-V SES-VFE	All	900	65,535	900	65,535	-	-
UAS-V UAS-VFE	All	900	65,535	900	65,535	-	-
FC-V FC-VFE	All	72	4,095	NR	NR	NR	NR
PPJC-VDet NPJC-VDet	All	32,767	2,097,151	32,767	2,097,151	-	-
PPJC-VGen NPJC-VGen	All	32,767	2,097,151	32,767	2,097,151	-	-
PJCDiff-V	All	32,767	2,097,151	32,767	2,097,151	-	-
PJCS-VDet PJCS-VGen	All	900	65,535	900	65,535	-	-

Notes:

- a: This value is based on the value that appeared in ANSI T1.231-1997, which was under study at the time that this document was released (see GR-253-ILR Issue ID 253-141).
- : To be determined
- NR: Not Required

6.2.2.2 Physical Layer PM

The intent of physical layer PM is to enable proactive monitoring of the physical devices and facilities that act as the transmitter, optical path and receiver of the SONET signal, so that an early indication of a problem is possible, before a failure actually occurs. These parameters are different from the other PM parameters in that the recorded value is a snapshot value recorded at a particular time during the current period, rather than a count that may change throughout the period. If the snapshot value is either less than or equal to a lower threshold value (if defined), or greater than or equal to an upper threshold value, then an out-of-range alert is sent to an OS. Each of the physical layer parameters defined here is a normalized value, referenced to a nominal value and expressed as a percentage. In some cases the nominal value is defined by the supplier, while in other cases it may be set

by the user (e.g., to be the expected value of the measurement when the device is first turned up).

6.2.2.2.1 Physical Layer Parameters

The physical layer performance parameters currently defined in SONET are shown below. Note that using these definitions, it is implied that the various nominal and measured values are expressed in linear units (e.g., mW or μ A) rather than logarithmic units (e.g., dBm). As discussed in GR-253-ILR Issue ID 253-141, these definitions and the possible use of logarithmic units are under study.

1. LBC_{normal} - This parameter is a measure of the Laser Bias Current (LBC). The normalized value of the LBC, expressed as an integer percentage, is the monitored parameter:

$$LBC_{normal} = \frac{LBC}{LBC_0} \times 100$$

where LBC_0 is the nominal value of the LBC provided by the NE supplier.

2. OPT_{normal} - This parameter is a measure of the average optical output power of the transmitter, or the Optical Power Transmitted (OPT). The normalized value of OPT, expressed as an integer percentage, is the monitored parameter:

$$OPT_{normal} = \frac{OPT}{OPT_0} \times 100$$

where OPT_0 is the nominal value of OPT provided by the NE supplier.

3. OPR_{normal} - This parameter is a measure of the average optical power of the received signal, or the Optical Power Received (OPR). The normalized value of OPR, expressed as an integer percentage, is the monitored parameter:

$$OPR_{normal} = \frac{OPR}{OPR_0} \times 100$$

where OPR_0 is the nominal value of OPR.

Unlike LBC_0 and OPT_0 , which are equipment-specific values that are not expected to need to be user-settable, OPR_0 is an installation or application-specific value. As such, its value needs to be user-settable. Three possible values to which it might be set are the receiver's specified maximum received power level (i.e., the overload power level) for installations with relatively little attenuation in the optical path, the receiver's specified minimum received power level (i.e., the receiver sensitivity) for installations where the optical path

attenuation is relatively large, or the power level received at equipment turn-up (assuming healthy equipment at turn-up).

6.2.2.2.2 Physical Layer PM Criteria

The following Physical layer PM criteria apply to all SONET NEs with optical interfaces.

- O6-311** [643v4] A SONET NE should support the LBC_{normal} and OPT_{normal} parameters to provide measurements of the health of each optical transmitter.
- O6-312** [994v3] A SONET NE should support the OPR_{normal} parameter to provide a measurement of the physical layer characteristics of the incoming signal at each optical receiver.
- R6-313** [1082] For each Physical layer PM parameter that it supports, the SONET NE shall measure and record the parameter value once per period. This snapshot value shall be recorded at approximately the same time (i.e., within ± 10 seconds) after the start of each new period.

To provide values for the user to retrieve in the current period, it is important for the NE to record its snapshot values as soon as practical after the start of each new period. For relatively small NEs, it may be practical to record the values "immediately" after the start of the new period, while for large NEs some delay may be necessary while other end-of-period/start-of period processing is performed. In any case, the following objective applies.

- O6-314** [1083] The SONET NE should record its Physical layer PM parameter snapshot values within one minute after the start of each new period.

In addition to snapshot values that a user might want to store and use for purposes such as long-term trend analyses, in some situations it might be desirable to be able to determine the current value of a Physical layer PM parameter. While continuous monitoring and retrieval capabilities for these parameters are not covered by the PM criteria in this section, they could be provided as separate Physical layer diagnostics (see Section 6.2.3.2.1). Alternatively, the NE could record a new snapshot value after a current period parameter register is initialized by the user. Although an invalid data flag would be set to indicate that the current period register had been initialized, this new snapshot value would actually be a valid value that would reflect the current value of the parameter at (approximately) the time that the register was initialized.

- O6-315** [1084] A SONET NE should record a new snapshot value within one minute after the current period register for a Physical layer PM parameter is initialized by the user.

As discussed in Section 6.2.2.1, an NE must generate an out-of-range alert when a snapshot value for a Physical layer PM parameter is not in its acceptable range. For the OPT_{normal} and OPR_{normal} parameters, the acceptable range has been defined to be based on both upper and lower thresholds. On the other hand, only an upper threshold is considered necessary for defining the acceptable range for the LBC_{normal} parameter.

- R6-316 [645v3]** A SONET NE shall provide lower threshold registers for the OPT_{normal} and OPR_{normal} parameters (if supported) and shall also provide an upper threshold register for each Physical layer PM parameter that it supports.

Similar to the case discussed above regarding nominal values, the acceptable ranges for the LBC_{normal} and OPT_{normal} parameters are generally expected to be equipment-specific. Therefore, the thresholds for those parameters are currently not required to be (but are allowed to be) user-settable. On the other hand, the acceptable range for the OPR_{normal} parameter is likely to be installation or application-specific. As such, that parameter's upper and lower thresholds need to be user-settable (i.e., the third bullet in **R6-304 [638v3]** is applicable for OPR_{normal}).

6.2.2.3 Section Layer PM

6.2.2.3.1 Section Layer Parameters

The following Section layer PM parameters are defined in SONET. (Note that no far-end PM parameters are defined for the Section layer.)

1. Section Severely Errored Framing Seconds (SEFS-Ss) – The SEFS-S parameter is a count of the seconds during which (at any point during the second) an SEF defect was present. See Section 5.5 for the definition of an SEF defect. (In addition, note that an SEF defect is expected to be present during most seconds in which an LOS or LOF defect is present. However, there may be situations when that is not the case, and the SEFS-S parameter is only incremented based on the presence of the SEF defect.)
2. Section Coding Violations (CV-Ss) – The CV-S parameter is a count of BIP errors detected at the Section layer (i.e., using the B1 byte in the incoming SONET signal). Up to eight Section BIP errors can be detected per STS-N frame, with each error incrementing the CV-S current second register.
3. Section Errored Seconds (ES-Ss) – The ES-S parameter is a count of the number of seconds during which (at any point during the second) at least one Section layer BIP error was detected or an SEF or LOS defect was present.

4. Section Severely Errored Seconds (SES-Ss) – The SES-S parameter is a count of the seconds during which K or more Section layer BIP errors were detected or an SEF or LOS defect was present. The number of BIP errors that cause a second to be considered an SES-S has been changed several times in both TA/TR/GR-253 and the applicable ANSI SONET standards, and thus may need to be settable.¹⁴ Table 6-9 contains the current values for K for the various SONET rates. These values are based on those that appear in ANSI T1.231-1997.

Table 6-9. Section BIP Errors to Trigger a Section SES

Rate	K
OC-1	52
OC-3	155
OC-12	616
OC-24	1,220
OC-48	2,392

6.2.2.3.2 Section Layer PM Criteria

The following requirements for section layer performance monitoring apply to all SONET NEs.

- R6-317** [656v2] A SONET NE shall be capable of accumulating the SEFS-S parameter.
- R6-318** [657] A SONET NE shall perform thresholding for the SEFS-S parameter.
- CR6-319** [658] A SONET NE may be required to support applications where line and section spans are not coincident.
- R6-320** [659v2] A SONET NE that supports applications where line and section spans are not coincident shall be capable of accumulating CV-Ss, ES-Ss, and SES-Ss.
- R6-321** [660v2] A SONET NE shall perform thresholding for the CV-S, ES-S and SES-S parameters if those parameters are supported.

14. The ability to set the value of K (at the Section, Line, STS Path, and VT Path layers) would be needed if future standards activities change the definitions of the SES parameters, and is not meant to imply that these numbers are variables.

Note that the B1 byte is not required to be generated for drop-side signals (i.e., an NE transmitting a drop-side signal may set the B1 byte to either the Section BIP-8 code or all-zeros), and that an NE that is capable of receiving a drop-side signal at a particular interface must be capable of ignoring the value in the incoming B1 byte (see Section 3.3.2.1). In addition, the value of the incoming B1 byte directly affects the accumulation of CV-Ss, ES-Ss, and SES-Ss.¹⁵ Therefore, if an NE supports only one or both of the following two Section PM accumulation options:

- accumulate all four of the defined Section PM parameters at a particular interface
- accumulate none of the defined Section PM parameters at a particular interface

(i.e., if it cannot be provisioned so that it just accumulates SEFS-Ss, or so that it ignores the B1 byte and bases its ES-S and SES-S parameters only on the presence of SEF and LOS defects), then it will not be capable either of receiving some drop-side signals, or of accumulating any valid Section PM data for those signals. Although the following objective is primarily intended to address this issue, the capability described in the first bullet item could also be useful in other applications with coincident section and line spans. Also note that if the capability described in the second bullet item is supported and used, the CV-S registers will not contain valid data (i.e., they will always contain a value of zero), and in most situations the ES-S, SES-S and SEFS-S registers will contain redundant data.

O6-322 [1033] A SONET NE that supports applications where line and section spans are not coincident should provide one of the following capabilities:

- A user-provisionable, per-section option to accumulate the CV-S, ES-S and SES-S parameters that is independent of any user-provisionable option to accumulate the SEFS-S parameter.
- A user-provisionable, per-section option to ignore the B1 byte for the purposes of accumulating the CV-S, ES-S and SES-S parameters.

6.2.2.4 Line Layer PM

6.2.2.4.1 Near-end Line Layer Parameters

The near-end (or incoming) Line layer PM parameters defined in SONET are:

1. Near-end Line Coding Violations (CV-Ls) – The CV-L parameter is a count of BIP errors detected at the Line layer (i.e., using the B2 bytes in the incoming

¹⁵ If one NE sets the B1 byte to all-zeros and another NE attempts to accumulate CV-Ss, ES-Ss, and SES-Ss using that value, none of those parameter registers will contain valid data. In such a situation, it is likely that every second will appear to contain enough BIP-8 errors to be considered a SES-S, causing CV-S accumulation to be inhibited, the continuous accumulation of ES-Ss and SES-Ss, and the generation of unnecessary TCAs.

SONET signal). Up to $8 \times N$ BIP errors can be detected per STS-N frame, with each error incrementing the CV-L current second register.

2. Near-end Line Errored Seconds (ES-Ls) – The ES-L parameter is a count of the seconds during which (at any point during the second) at least one Line layer BIP error was detected or an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) was present.
3. Near-end Line Severely Errored Seconds (SES-Ls) – The SES-L parameter is a count of the seconds during which K or more Line layer BIP errors were detected or an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) was present. The number of BIP errors that cause a second to be considered an SES-L has been changed several times in both TA/TR/GR-253 and the applicable ANSI SONET standards, and may need to be settable. Table 6-10 contains the current values for K for the various SONET rates. These values are based on those that appear in ANSI T1.231-1997.

**Table 6-10. Line BIP Errors to Trigger a Line SES
(Near-End and Far-End)**

Rate	K
OC-1	51
OC-3	154
OC-12	615
OC-24	1,230
OC-48	2,459

4. Near-end Line Unavailable Seconds (UAS-Ls) – The UAS-L parameter is a count of the seconds during which the Line was considered unavailable. A Line becomes unavailable at the onset of 10 consecutive seconds that qualify as SES-Ls, and continues to be unavailable until the onset of 10 consecutive seconds that do not qualify as SES-Ls.
5. Near-end Line Failure Counts (FC-Ls) – The FC-L parameter is a count of the number of near-end line failure events. A failure event begins when the AIS-L failure (or a lower-layer, traffic-related, near-end failure, see Section 6.2.1.8.2) is declared, and ends when the failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins. Note that functionally, an AIS-L failure will be declared either by receiving (and timing) an AIS-L signal from another NE, or by receiving (and timing) an internally generated AIS-L signal from STE in the same NE where the LTE resides.

6. Protection Switching Counts (PSCs) – For a working line, the PSC parameter is a count of the number of times that service has been switched from the monitored line to the protection line, plus the number of times it has been switched back to the working line. For the protection line, it is a count of the number of times that service has been switched from any working line to the protection line, plus the number of times service has been switched back to a working line. The PSC parameter is only applicable if line level protection switching is used.
7. Protection Switching Duration (PSD) – For a working line, the PSD parameter is a count of the seconds that service was being carried on the protection line. For the protection line, it is a count of the seconds that the line was being used to carry service. The PSD parameter is only applicable if revertive line level protection switching is used.

6.2.2.4.2 Far-end Line Layer Parameters

Far-end Line layer performance is conveyed back to the near-end LTE via the K2 byte (RDI-L) and the M0 or M1 byte (REI-L). The far-end Line layer PM parameters defined in SONET are:

1. Far-end Line Coding Violations (CV-LFEs) – The CV-LFE parameter is a count of the number of BIP errors detected by the far-end LTE and reported back to the near-end LTE using the REI-L indication in the Line overhead. For SONET signals at rates below OC-48, up to $8 \times N$ BIP errors per STS-N frame can be indicated using the REI-L. For OC-48 signals, up to 255 BIP errors per STS-N frame can be indicated. The CV-LFE current second register is incremented for each BIP error indicated by the incoming REI-L.
2. Far-end Line Errored Seconds (ES-LFEs) – The ES-LFE parameter is a count of the seconds during which (at any point during the second) at least one Line BIP error was reported by the far-end LTE (using the REI-L indication) or an RDI-L defect was present.
3. Far-end Line Severely Errored Seconds (SES-LFEs) – The SES-LFE parameter is a count of the seconds during which K or more Line BIP errors were reported by the far-end LTE or an RDI-L defect was present. The number of reported far-end BIP errors that cause a second to be considered an SES-LFE may need to be settable. Table 6-10 contains the current values for K for the various SONET rates.
4. Far-end Line Unavailable Seconds (UAS-LFE) – The UAS-LFE parameter is a count of the seconds during which the Line is considered unavailable at the far end. A Line is considered unavailable at the far end at the onset of 10 consecutive seconds that qualify as SES-LFEs, and continues to be considered unavailable until the onset of 10 consecutive seconds that do not qualify as SES-LFEs.

5. Far-end Line Failure Counts (FC-LFEs) – The FC-LFE parameter is a count of the number of far-end line failure events. A failure event begins when the RFI-L failure is declared, and ends when the RFI-L failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins.

6.2.2.4.3 Line Layer PM Criteria

The following requirements for Line layer performance monitoring apply to all SONET NEs that terminate the Line layer.

- R6-323** [661] A SONET NE providing LTE functions shall accumulate CV-Ls, ES-Ls, SES-Ls, UAS-Ls, and FC-Ls for each line.
- R6-324** [662] A SONET NE providing LTE functions shall perform thresholding for the CV-L, ES-L, SES-L, and UAS-L parameters.
- R6-325** [663] A SONET NE that supports line protection switching for a given line shall accumulate PSCs for that line.
- R6-326** [664] A SONET NE that is using revertive protection switching for a given line shall accumulate PSDs for that line.

| Thresholding is not required for the PSC, PSD, and FC-L parameters.

- R6-327** [667] A SONET NE providing LTE functions shall provide the capability, on a per-line basis, to accumulate the CV-LFE, ES-LFE, SES-LFE, UAS-LFE, and FC-LFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be “not active.”
- R6-328** [668] A SONET NE that is accumulating far-end line PM parameters shall perform thresholding for the CV-LFE, ES-LFE, SES-LFE, and UAS-LFE parameters.

6.2.2.5 STS Path Layer PM

6.2.2.5.1 Near-end STS Path Layer Parameters

The near-end STS Path layer PM parameters defined in SONET are:

1. Near-end STS Path Coding Violations (CV-Ps) – The CV-P parameter is a count of BIP errors detected at the STS Path layer (i.e., using the B3 byte in the incoming STS Path overhead). Up to 8 BIP errors can be detected per frame, with each error incrementing the CV-P current second register.
2. Near-end STS Path Errored Seconds (ES-Ps) – The ES-P parameter is a count of the seconds during which (at any point during the second) at least one STS Path BIP error was detected, or an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-P defect or, if the STS PTE monitoring the path supports ERDI-P for that path, an UNEQ-P or TIM-P defect was present.
3. Near-end STS Path Severely Errored Seconds (SES-Ps) – The SES-P parameter is a count of the seconds during which K or more STS Path BIP errors were detected, or an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-P defect or, if the STS PTE monitoring the path supports ERDI-P for that path, an UNEQ-P or TIM-P defect was present. The number of BIP errors that cause a second to be considered an SES-P may need to be settable. As shown in Table 6-11, the current values for K (which are based on the those that appear in ANSI T1.231) are independent of the particular size of the STS path (e.g., STS-1, STS-3c).

**Table 6-11. STS Path BIP Errors to Trigger an STS Path SES
(Near-End and Far-End)**

Rate	K
STS-1	2,400
STS-3c	2,400
STS-12c	2,400
STS-48c	2,400

4. Near-end STS Path Unavailable Seconds (UAS-Ps) – The UAS-P parameter is a count of the seconds during which the STS Path was considered unavailable. An STS Path becomes unavailable at the onset of 10 consecutive seconds that qualify as SES-Ps, and continues to be unavailable until the onset of 10 consecutive seconds that do not qualify as SES-Ps.
5. Near-end STS Path Failure Counts (FC-P) – The FC-P parameter is a count of the number of near-end STS Path failure events. A failure event begins when an AIS-P failure (or a lower-layer, traffic-related, near-end failure, see Section 6.2.1.8.2), an LOP-P failure or, if the STS PTE monitoring the path supports ERDI-P for that path, an UNEQ-P or TIM-P failure is declared. The failure event ends when these failures are cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins. Note that functionally, an

AIS-P failure will be declared either by receiving (and timing) an AIS-P signal from another NE, or by receiving (and timing) an internally generated AIS-P signal from LTE in the same NE where the STS PTE resides.

6. Positive Pointer Justification Count - STS Path Detected (PPJC-PDet) – The PPJC-PDet parameter is a count of the positive pointer justifications (i.e., valid increment operations) detected on a particular path in an incoming SONET signal.
7. Negative Pointer Justification Count - STS Path Detected (NPJC-PDet) – The NPJC-PDet parameter is a count of the negative pointer justifications (i.e., valid decrement operations) detected on a particular path in an incoming SONET signal.
8. Positive Pointer Justification Count - STS Path Generated (PPJC-PGen) – The PPJC-PGen parameter is a count of the positive pointer justifications (i.e., increment operations) generated for a particular path to reconcile the frequency of the SPE with the local clock.
9. Negative Pointer Justification Count - STS Path Generated (NPJC-PGen) – The NPJC-PGen parameter is a count of the negative pointer justifications (i.e., decrement operations) generated for a particular path to reconcile the frequency of the SPE with the local clock.
10. Pointer Justification Count Difference - STS Path (PJCDiff-P) – The PJCDiff-P parameter is the absolute value of the difference between the net number of detected pointer justification counts and the net number of generated pointer justification counts. That is, PJCDiff-P is equal to $|(PPJC-PGen - NPJC-PGen) - (PPJC-PDet - NPJC-PDet)|$.
11. Pointer Justification Count Seconds - STS Path Detect (PJCS-PDet) – The PJCS-PDet parameter is a count of the one-second intervals containing one or more PPJC-PDet or NPJC-PDet.
12. Pointer Justification Count Seconds - STS Path Generate (PJCS-PGen) – The PJCS-PGen parameter is a count of the one-second intervals containing one or more PPJC-PGen or NPJC-PGen.

Note that parameters 6 through 12 are sometimes referred to collectively as the “STS PJ-related” parameters.

6.2.2.5.2 Far-end STS Path Layer Parameters

Far-end STS Path layer performance is conveyed back to the near-end STS PTE via bits 1 through 4 (REI-P) and 5 through 7 (RDI-P) of the G1 byte. The far-end STS Path layer PM parameters defined in SONET are:

1. Far-end STS Path Coding Violations (CV-PFEs) – The CV-PFE parameter is a count of the number of BIP errors detected by the far-end STS PTE and reported

back to the near-end STS PTE using the REI-P indication in the STS Path overhead. Up to 8 BIP errors per frame can be indicated. The CV-PFE current second register is incremented for each BIP error indicated by the incoming REI-P.

2. Far-end STS Path Errored Seconds (ES-PFEs) – The ES-PFE parameter is a count of the seconds during which (at any point during the second) at least one STS Path BIP error was reported by the far-end STS PTE (using the REI-P indication), a one-bit RDI-P defect was present, or (if ERDI-P is supported, see Section 6.2.1.3.2) an ERDI-P Server or Connectivity defect was present.
3. Far-end STS Path Severely Errored Seconds (SES-PFEs) – The SES-PFE parameter is a count of the seconds during which K or more STS Path BIP errors were reported by the far-end STS PTE, a one-bit RDI-P defect was present, or (if ERDI-P is supported) an ERDI-P Server or Connectivity defect was present. The number of reported far-end BIP errors that cause a second to be considered an SES-PFE may need to be settable. Table 6-11 contains the current values for K for STS paths.
4. Far-end STS Path Unavailable Seconds (UAS-PFE) – The UAS-PFE parameter is a count of the seconds during which the STS Path is considered unavailable at the far end. An STS Path is considered unavailable at the far end at the onset of 10 consecutive seconds that qualify as SES-PFEs, and continues to be considered unavailable until the onset of 10 consecutive seconds that do not qualify as SES-PFEs.
5. Far-end STS Path Failure Counts (FC-PFEs) – The FC-PFE parameter is a count of the number of far-end STS Path failure events. A failure event begins when a one-bit RFI-P failure, or (if ERDI-P is supported) an ERFI-P Server or Connectivity failure is declared. The failure event ends when the RFI-P failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins.

Note that with the parameter definitions shown above, unless the STS PTE at both ends of a path support the same version of RDI-P, inconsistencies can be expected to occur between the near-end PM data accumulated at one NE and the far-end data accumulated at the far-end NE. The reason for this is that connectivity defects and failures (e.g., UNEQ-P, ERDI-P Connectivity) are included in the definitions of the near-end and far-end ES-P, SES-P and FC-P parameters if ERDI-P is supported, but are not included if one-bit RDI-P is being used.

6.2.2.5.3 STS Path Layer PM Criteria

The following requirements for STS Path layer performance monitoring apply to all SONET NEs that terminate the STS Path layer, except the SONET digital switch trunk side interface (see TR-TSY-000782, *SONET Digital Switch Trunk Interface Criteria*).

R6-329 [669] A SONET NE providing STS PTE functions shall accumulate CV-Ps, ES-Ps, SES-Ps, UAS-Ps, and FC-Ps for each terminated STS Path.

R6-330 [670] A SONET NE providing STS PTE functions shall perform thresholding for the CV-P, ES-P, SES-P, and UAS-P parameters.

It is not necessary for a SONET NE to accumulate VT PJs if all VT Paths are terminated. A SONET NE may optionally accumulate VT PJs on a nonterminating VT SPE received from each NE at which STS PJs are not accumulated on any nonterminated STS Path.

CR6-331 [1085] STS PTE that processes the STS pointer from an incoming SONET signal (i.e., STS PTE that terminates a path that has not been reconciled to the local clock by upstream LTE within the same NE) may be required to support the capability to accumulate the STS PJ-related parameters defined in Section 6.2.2.5.1.

Note that although they are included for consistency with the case where the STS PJ-related parameters are accumulated (as intermediate-path PM parameters) at LTE, the PPJC-PGen, NPJC-PGen and PJCS-PGen parameters are expected to always be equal to zero when they are accumulated at STS PTE.

R6-332 [1086] If the accumulation of the STS PJ-related parameters is supported, the user shall be able to activate that accumulation on a per-path basis. The default setting shall be "not active".

R6-333 [1087] A SONET NE that is accumulating STS PJ-related parameters shall perform thresholding for the PPJC-PDet, NPJC-PDet, PPJC-PGen, NPJC-PGen, PJCDiff-P, PJCS-PDet and PJCS-PGen parameters.

Note that unlike all of the other PM parameters that are based on integer counts, the PJCDiff-P and (at the VT layer) PJCDiff-V parameters are not essentially monotonically increasing functions during any particular accumulation period. Therefore it is possible that their values may reach or exceed the corresponding thresholds for generating TCAs, decrease to values less than the thresholds, and then increase back to the thresholds, etc. Existing requirements indicate that multiple TCAs must not be generated in such situations (see **R820-38** and **R820-39** in GR-820-CORE); however, it may still be desirable to reduce the chance that those situations will occur. This can be done by providing default threshold values that are large enough that they will only be reached if there is frequency offset in the system (see GR-253-ILR Issue ID 253-30).

R6-334 [672] A SONET NE providing STS PTE functions shall provide the capability, on a per STS Path basis, to accumulate the CV-PFE, ES-PFE, SES-PFE, UAS-PFE, and FC-PFE parameters, and to activate and

deactivate the accumulation of these parameters (as a group). The default setting shall be "not active."

- R6-335** [673] A SONET NE that is accumulating far-end STS Path PM parameters shall perform thresholding for the CV-PFE, ES-PFE, SES-PFE, and UAS-PFE parameters.

6.2.2.6 VT Path Layer PM

6.2.2.6.1 Near-end VT Path Layer Parameters

The near-end VT Path layer PM parameters defined in SONET are:

1. Near-end VT Path Coding Violations (CV-Vs) – The CV-V parameter is a count of BIP errors detected at the VT Path layer (i.e., using bits 1 and 2 of the V5 byte in the incoming VT Path overhead). Up to 2 BIP errors can be detected per VT superframe, with each error incrementing the CV-V current second register.
2. Near-end VT Path Errored Seconds (ES-Vs) – The ES-V parameter is a count of the seconds during which (at any point during the second) at least one VT Path BIP error was detected, or an AIS-V defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-V defect or, if the VT PTE monitoring the path supports ERDI-V for that path, an UNEQ-V defect was present.
3. Near-end VT Path Severely Errored Seconds (SES-Vs) – The SES-V parameter is a count of the seconds during which K or more VT Path BIP errors were detected, or an AIS-V defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-V defect or, if the VT PTE monitoring the path supports ERDI-V for that path, an UNEQ-V defect was present. The number of BIP errors that cause a second to be considered an SES-V may need to be settable. As shown in Table 6-12, the current values for K (which are based on the those that appear in ANSI T1.231) are independent of the particular size of the VT path (e.g., VT1.5, VT2).

**Table 6-12. VT Path BIP Errors to Trigger a VT Path SES
(Near-End and Far-End)**

Rate	K
VT1.5	600
VT2	600
VT3	600
VT6	600

4. Near-end VT Path Unavailable Seconds (UAS-Vs) – The UAS-V parameter is a count of the seconds during which the VT Path was considered unavailable. A VT Path becomes unavailable at the onset of 10 consecutive seconds that qualify as SES-Vs, and continues to be unavailable until the onset of 10 consecutive seconds that do not qualify as SES-Vs.
5. Near-end VT Path Failure Counts (FC-V) – The FC-V parameter is a count of the number of near-end VT Path failure events. A failure event begins when an AIS-V failure (or a lower-layer, traffic-related, near-end failure, see Section 6.2.1.8.2), an LOP-V failure or, if the VT PTE monitoring the path supports ERDI-V for that path, an UNEQ-V failure is declared. The failure event ends when these failures are cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins. Note that functionally, an AIS-V failure will be declared either by receiving (and timing) an AIS-V signal from another NE, or by receiving (and timing) an internally generated AIS-V signal from STS PTE in the same NE where the VT PTE resides.
6. Positive Pointer Justification Count - VT Path Detected (PPJC-VDet) – The PPJC-VDet parameter is a count of the positive pointer justifications (i.e., valid increment operations) detected on a particular path in an incoming SONET signal.
7. Negative Pointer Justification Count - VT Path Detected (NPJC-VDet) – The NPJC-VDet parameter is a count of the negative pointer justifications (i.e., valid decrement operations) detected on a particular path in an incoming SONET signal.
8. Positive Pointer Justification Count - VT Path Generated (PPJC-VGen) – The PPJC-VGen parameter is a count of the positive pointer justifications (i.e., increment operations) generated for a particular path to reconcile the frequency of the SPE with the local clock.
9. Negative Pointer Justification Count - VT Path Generated (NPJC-VGen) – The NPJC-VGen parameter is a count of the negative pointer justifications (i.e., decrement operations) generated for a particular path to reconcile the frequency of the SPE with the local clock.

10. Pointer Justification Count Difference - VT Path (PJCDiff-V) – The PJCDiff-V parameter is the absolute value of the difference between the net number of detected pointer justification counts and the net number of generated pointer justification counts. That is, PJCDiff-V is equal to $|(PPJC-VGen - NPJC-VGen) - (PPJC-VDet - NPJC-VDet)|$.
11. Pointer Justification Count Seconds - VT Path Detect (PJCS-VDet) – The PJCS-VDet parameter is a count of the one-second intervals containing one or more PPJC-VDet or NPJC-VDet.
12. Pointer Justification Count Seconds - VT Path Generate (PJCS-VGen) – The PJCS-VGen parameter is a count of the one-second intervals containing one or more PPJC-VGen or NPJC-VGen.

Note that parameters 6 through 12 are sometimes referred to collectively as the “VT PJ-related” parameters.

6.2.2.6.2 Far-end VT Path Layer Parameters

Far-end VT Path layer performance is conveyed back to the near-end VT PTE via bit 3 of the V5 byte (REI-V), and either bits 5 through 7 of the Z7 byte or bit 8 of the V5 byte (RDI-V). The far-end T Path layer PM parameters defined in SONET are:

1. Far-end VT Path Coding Violations (CV-VFEs) – The CV-VFE parameter is a count of the number of BIP errors detected by the far-end VT PTE and reported back to the near-end VT PTE using the REI-V indication in the VT Path overhead. Note that only 1 BIP error can be indicated per VT superframe using the REI-V bit (out of the two BIP errors that can be detected). The CV-VFE current second register is incremented for each BIP error indicated by the incoming REI-V.
2. Far-end VT Path Errored Seconds (ES-VFEs) – The ES-VFE parameter is a count of the seconds during which (at any point during the second) at least one VT Path BIP error was reported by the far-end VT PTE (using the REI-V indication), a one-bit RDI-V defect was present, or (if ERDI-V is supported, see Section 6.2.1.3.3) an ERDI-V Server or Connectivity defect was present.
3. Far-end VT Path Severely Errored Seconds (SES-VFEs) – The SES-VFE parameter is a count of the seconds during which K or more VT Path BIP errors were reported by the far-end VT PTE, a one-bit RDI-V defect was present, or (if ERDI-V is supported) an ERDI-V Server or Connectivity defect was present. The number of reported far-end BIP errors that cause a second to be considered an SES-VFE may need to be settable. Table 6-12 contains the current values for K for the various VT paths.
4. Far-end VT Path Unavailable Seconds (UAS-VFE) – The UAS-VFE parameter is a count of the seconds during which the VT Path is considered unavailable at the

far end. An VT Path is considered unavailable at the far end at the onset of 10 consecutive seconds that qualify as SES-VFEs, and continues to be considered unavailable until the onset of 10 consecutive seconds that do not qualify as SES-VFEs.

5. Far-end VT Path Failure Counts (FC-VFEs) – The FC-VFE parameter is a count of the number of far-end VT Path failure events. A failure event begins when a one-bit RFI-V failure, or (if ERDI-V is supported) an ERFI-V Server or Connectivity failure is declared. The failure event ends when the RFI-V failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins.

Note that with the parameter definitions shown above, unless the VT PTE at both ends of a path support the same version of RDI-V, inconsistencies can be expected to occur between the near-end PM data accumulated at one NE and the far-end data accumulated at the far-end NE. The reason for this is that connectivity defects and failures (e.g., UNEQ-V, ERDI-V Connectivity) are included in the definitions of the near-end and far-end ES-V, SES-V and FC-V parameters if ERDI-V is supported, but are not included if one-bit RDI-V is being used.

6.2.2.6.3 VT Path Layer PM Criteria

The following requirements for VT Path layer performance monitoring apply to all SONET NEs that terminate the VT Path layer.

- | | |
|----------------|--|
| R6-336 | [674] A SONET NE providing VT PTE functions shall accumulate CV-Vs, ES-Vs, SES-Vs, UAS-Vs, and FC-Vs for each terminated VT Path. |
| R6-337 | [675] A SONET providing VT PTE functions shall provide thresholding for the CV-V, ES-V, SES-V, and UAS-V parameters. |
| CR6-338 | [1088] VT PTE that processes the VT pointer from an incoming SONET signal (i.e., VT PTE that terminates a path that has not been reconciled to the local clock by upstream STS PTE within the same NE) may be required to support the capability to accumulate the VT PJ-related parameters defined in Section 6.2.2.6.1. |

Note that although they are included for consistency with the case where the VT PJ-related parameters are accumulated (as intermediate-path PM parameters) at STS PTE, the PPJC-VGen, NPJC-VGen and PJCS-VGen parameters are expected to always be equal to zero when they are accumulated at VT PTE.

- R6-339** [1089] If the accumulation of the VT PJ-related parameters is supported, the user shall be able to activate that accumulation on a per-path basis. The default setting shall be "not active".
- R6-340** [1090] A SONET NE that is accumulating VT PJ-related parameters shall perform thresholding for the PPJC-VDet, NPJC-VDet, PPJC-VGen, NPJC-VGen, PJCDiff-V, PJCS-VDet and PJCS-VGen parameters.
- R6-341** [676] A SONET NE providing VT PTE functions shall provide the capability to accumulate, on a per VT Path basis, the CV-VFE, ES-VFE, SES-VFE, UAS-VFE, and FC-VFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be "not active."
- R6-342** [677] A SONET NE that is accumulating the Far-end VT Path PM parameters shall perform thresholding for the CV-VFE, ES-VFE, SES-VFE, and UAS-VFE parameters.

6.2.2.7 Monitoring at DS_n Interfaces

GR-820-CORE defines DS_n path and line PM parameters.

- R6-343** [678] A SONET NE shall provide DS_n path PM for each DS_n path that it terminates.
- R6-344** [679] A SONET NE shall provide DS_n line PM for each DS_n line that it terminates.
- R6-345** [680] A SONET NE shall meet the accumulation and thresholding requirements in GR-820-CORE for DS_n Line and Path PM parameters.
- R6-346** [681] A SONET NE shall meet the OS reporting and data retrieval requirements in GR-820-CORE for its DS_n PM parameters.

Note that the accumulation of DS_n path PM where the DS_n path is not terminated is considered to be intermediate-path PM, which is covered in Section 6.2.2.9.

6.2.2.8 PM During Troubles

The criteria concerning PM during troubles depend on the entity being monitored and the particular PM parameter. In general, the accumulation of a parameter is either inhibited during unavailable time, inhibited during unavailable time and during severely errored

seconds, or never inhibited. For example, the accumulation of the ES and SES parameters is inhibited during periods of unavailability of the monitored entity. This is accomplished by not incrementing the current 15-minute and current-day registers during unavailable seconds, where unavailable seconds are defined in Sections 6.2.2.4, 6.2.2.5, and 6.2.2.6 (for SONET Lines, STS Paths and VT Paths). The accumulation of CVs is also inhibited during unavailable time; however, it is also inhibited for all seconds that are considered to be SESs for that entity (i.e., seconds during which certain defects are present, or during which K or more CVs are detected).

The accumulation of several PM parameters is never inhibited. For example, the accumulation of UASs and FCs (where they are defined) is not inhibited because they keep track of unavailable time and failure counts for the monitored entity. In addition, the accumulation of several of the SONET Section layer PM parameters is never inhibited because a UAS parameter is not defined at that layer.

Note that with these definitions, a CV positive adjustment register is still needed to keep track of the number of CVs that occur (at a sub-SES rate) during the 10 seconds it takes to get out of unavailable time. In addition, those 10 seconds could easily span two periods so the NE may need to keep track of how many of the CVs occurred in each period, and possibly send a TCA for the previous period if the adjustment puts that period over the threshold (see Section 6.2.2.1).

R6-347 [683v3] A SONET NE shall inhibit the accumulation of near-end SONET PM parameters according to the rules listed below (and as summarized in Tables 6-13 through 6-19).

- ... • The accumulation of all near-end Line layer parameters except for UASs, FCs, PSCs, and PSDs (if applicable) shall be inhibited during periods of unavailability of the monitored line.
- ... • The accumulation of all near-end STS or VT layer parameters except for UASs and FCs shall be inhibited during periods of unavailability of the monitored path.
- ... • The accumulation of near-end CVs for a particular entity shall be inhibited for all seconds that are counted as SESs for that entity.

Note that PJ-related parameters are not listed as exceptions in the second bullet of **R6-347 [683v3]**, and therefore the accumulation of those parameters for a particular STS or VT Path is currently required to be inhibited during periods of unavailability for that path. This is consistent with the specifications contained in ANSI T1.231-1997 and appears to be feasible in certain situations (i.e., when the parameters are being accumulated at PTE terminating that path, or when both the PJ-related and non-PJ-related parameters are being accumulated for a level-1 intermediate path). However, in certain other situations (i.e., when only the PJ-related parameters are being accumulated for a level-1 intermediate path) it may be desirable for them to be accumulated continuously, making it unnecessary for the NE to monitor the path overhead and detect various defects. This issue is for further study.

Far-end PM parameters are derived from information carried in the incoming signal, and therefore can only be properly accumulated when that information is available. Therefore, the accumulation of all far-end parameters is also inhibited for seconds during which certain near-end defects are present.

R6-348 [995v2] A SONET NE shall inhibit the accumulation of far-end SONET PM parameters according to the rules listed below (and as summarized in Tables 6-13 through 6-19).

- ... • The accumulation of all far-end parameters except for far-end UASs and far-end FCs shall be inhibited during periods of unavailability of the line or path at the far end.
- ... • The accumulation of far-end CVs for a particular entity shall be inhibited for all seconds that are counted as (far-end) SESs for that entity.
- ... • The accumulation of all far-end Line parameters shall be inhibited during seconds in which a near-end AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) is present. An invalid-data flag shall be set for the inhibited parameters.
- ... • The accumulation of all far-end STS Path parameters shall be inhibited during seconds in which a near-end AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), a near-end LOP-P defect, or a near-end UNEQ-P defect is present. An invalid-data flag shall be set for the inhibited parameters.
- ... • The accumulation of all far-end VT Path parameters shall be inhibited during seconds in which a near-end AIS-V defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), a near-end LOP-V defect, or a near-end UNEQ-V defect is present. An invalid-data flag shall be set for the inhibited parameters.

R6-349 [682] A SONET NE that provides DS_n performance monitoring shall meet the requirements in GR-820-CORE for DS_n PM during troubles.

The following abbreviations are used in Tables 6-13 through 6-19 to indicate if the accumulation of a parameter continues when the listed defect is present and during unavailable time.

- y - Indicates the parameter shall continue to be accumulated during seconds in which the defect is present.
- N - Indicates the accumulation of the parameter shall be inhibited for all seconds during which the defect is present, as well as during periods of unavailability.

- n - Indicates the parameter shall be accumulated during seconds during which the defect is present, until unavailable time is declared for the given entity. When unavailable time is declared, the accumulation of the parameter shall be inhibited, back to the point when unavailable time began.
- 0 - Indicates the accumulation of the (far-end) parameter shall be inhibited for each second during which the (near-end) defect is present on the incoming signal. In such cases, the parameter shall be marked as invalid for the entire accumulation interval.

Note that only defects that are detected at, or below, the layer where a parameter is accumulated can affect the accumulation of that parameters. For example, the accumulation of Section layer PM parameters is not affected by an AIS-L defect, which is detected at the Line layer. Therefore, AIS-L is not listed in Table 6-13 on Section layer PM accumulation, but is listed in Tables 6-14 through 6-19 on Line, STS Path, and VT Path layer PM accumulation. Finally, RDI defects do not affect the accumulation of near-end PM parameters, and therefore are not listed in Tables 6-13 through 6-16.

Table 6-13. Section Layer PM Accumulation During Defects

(Near-End) Section Parameter	Defect ^a
	LOS or SEF
SEFS-S	y ^b
CV-S	N
ES-S	y
SES-S	y

Notes:

- a. Defects detected at the Line and Path layers do not affect the accumulation of Section layer PM.
- b. See page 6-123 for definitions of y and N.

Table 6-14. Near-End Line Layer PM Accumulation During Defects

Near-End Line Parameter	Defect ^a	
	LOS, LOF, or AIS-L ^b	LOP-P, or all-ones STS pointer relay ^c
CV-L	N ^d	y
ES-L	n	y
SES-L	n	y
UAS-L	y	y
FC-L	y	y
PSC	y	y
PSD	y	y

Notes:

- RDI-L defects and defects detected at the Path layers do not affect the accumulation of near-end Line layer PM.
- In the functional model used in this document, LOS and LOF defects are detected at STE, which then generates AIS-L downstream. This AIS-L is detected at the LTE and affects the accumulation of Line layer PM.
- Although LTE that processes STS pointers detects LOP-P defects and performs all-ones STS pointer relay, neither of those processes affect the accumulation of near-end Line PM parameters.
- See page 6-123 for definitions of y, N, and n.

Table 6-15. Near-End STS Path Layer PM Accumulation During Defects

Near-End STS Path Parameter	Defect ^a		
	LOS, LOF, AIS-L, AIS-P ^b , LOP-P, UNEQ-P ^c or TIM-P ^c	UNEQ-P ^c , TIM-P ^c , or PLM-P	LOP-V, or all-ones VT pointer relay ^d
CV-P	N ^e	y	y
ES-P	n	y	y
SES-P	n	y	y
UAS-P	y	y	y
FC-P	y	y	y
STS PJ-related parameters	n ^f	y	y

Notes:

- RDI-L and RDI-P defects, and defects detected at the VT Path layer do not affect the accumulation of near-end STS Path layer PM.
- In the functional model used in this document, defects detected at STE or LTE (e.g., LOS) result in AIS-P being generated downstream. This AIS-P is detected at the STS PTE and affects the accumulation of STS Path layer PM.
- If the STS PTE monitoring the path supports ERDI-P for that path, then UNEQ-P and TIM-P (if activated) are included in the second column of this table (along with LOS, LOF, etc.). On the other hand, if one-bit RDI-P is supported, then UNEQ-P and TIM-P are included in the third column (with PLM-P).
- Although STS PTE that processes VT pointers detects LOP-V defects and performs all-ones VT pointer relay, neither of those processes affect the accumulation of near-end STS Path PM parameters.
- See page 6-123 for definitions of y, N, and n.
- The criteria on STS PJ-related parameter accumulation during troubles are under study (see page 6-122).

Table 6-16. Near-End VT Path Layer PM Accumulation During Defects

Near-End VT Path Parameter	Defect ^a	
	LOS, LOF, AIS-L, LOP-P, AIS-P, UNEQ-P, TIM-P, PLM-P, AIS-V ^b , LOP-V, or UNEQ-V ^c	UNEQ-V ^c , or PLM-V
CV-V	N ^d	y
ES-V	n	y
SES-V	n	y
UAS-V	y	y
FC-V	y	y
VT PJ-related parameters	n ^e	y

Notes:

- RDI defects do not affect the accumulation of near-end VT Path layer PM.
- In the functional model used in this document, defects detected at STE, LTE, or STS PTE (e.g., LOS) result in AIS-V being generated downstream. This AIS-V is detected at the VT PTE and affects the accumulation of VT Path layer PM.
- If the VT PTE monitoring the path supports ERDI-V for that path, then UNEQ-V is included in the second column of this table (along with LOS, LOF, etc.). On the other hand, if one-bit RDI-V is supported, then UNEQ-V is included in the third column (with PLM-V).
- See page 6-123 for definitions of y, N, and n.
- The criteria on VT PJ-related parameter accumulation during troubles are under study (see page 6-122).

Table 6-17. Far-End Line Layer PM Accumulation During Defects

Far-End Line Parameter	Defect ^a		
	LOS, LOF, or AIS-L ^b	RDI-L	LOP-P, or all-ones STS pointer relay ^c
CV-LFE	0 ^d	N	y
ES-LFE	0	n	y
SES-LFE	0	n	y
UAS-LFE	0	y	y
FC-LFE	0	y	y

Notes:

- Defects detected at the Path layers do not affect the accumulation of far-end Line layer PM.
- In the functional model used in this document, LOS and LOF defects are detected at STE, which then generates AIS-L downstream. This AIS-L is detected at the LTE and affects the accumulation of Line layer PM.
- Note that although all-ones STS pointer relay is performed by LTE that processes STS pointers, it is not defined as a defect and does not affect the accumulation of far-end Line PM parameters.
- See page 6-123 for definitions of y, N, n, and 0.

Table 6-18. Far-End STS Path Layer PM Accumulation During Defects

Far-End STS Path Parameter	Defect ^a				
	LOS, LOF, AIS-L, AIS-P ^b , LOP-P, or UNEQ-P	RDI-L, TIM-P, or PLM-P	One-bit RDI-P, ERDI-P Server, or ERDI-P Connectivity ^c	ERDI-P Payload	LOP-V, or all-ones VT pointer relay ^d
CV-PFE	0 ^c	y	N	y	y
ES-PFE	0	y	n	y	y
SES-PFE	0	y	n	y	y
UAS-PFE	0	y	y	y	y
FC-PFE	0	y	y	y	y

Notes:

- Defects detected at the VT Path layer do not affect the accumulation of far-end STS Path layer PM.
- In the functional model used in this document, defects detected at STE or LTE (e.g., LOS) result in AIS-P being generated downstream. This AIS-P is detected at the STS PTE and affects the accumulation of STS Path layer PM.
- See Section 6.2.1.3.2 for information on the various RDI-P defects. Also note that the effect that an UNEQ-P or TIM-P defect detected by the far-end STS PTE will have on the accumulation of far-end STS Path layer PM parameters (by the near-end STS PTE) depends on whether the near-end NE supports one-bit or enhanced RDI-P, and also on which RDI-P is supported by the far-end NE.
- Note that although all-ones VT pointer relay is performed by STS PTE that processes VT pointers, it is not defined as a defect and does not affect the accumulation of far-end STS Path PM parameters.
- See page 6-123 for definitions of y, N, n, and 0.

Table 6-19. Far-End VT Path Layer PM Accumulation During Defects

Far-End VT Path Parameter	Defect			
	LOS, LOF, AIS-L, LOP-P, AIS-P, UNEQ-P, TIM-P, PLM-P, AIS-V ^a , LOP-V, or UNEQ-V	RDI-L, RDI-P, or PLM-V	One-bit RDI-V, ERDI-V Server, or ERDI-V Connectivity ^b	ERDI-V Payload
CV-VFE	0 ^c	y	N	y
ES-VFE	0	y	n	y
SES-VFE	0	y	n	y
UAS-VFE	0	y	y	y
FC-VFE	0	y	y	y

Notes:

- In the functional model used in this document, defects detected at STE, LTE, or STS PTE (e.g., LOS) result in AIS-V being generated downstream. This AIS-V is detected at the VT PTE and affects the accumulation of VT Path layer PM.
- See Section 6.2.1.3.3 for information on the various RDI-V defects. Also note that the effect that an UNEQ-V defect detected by the far-end VT PTE will have on the accumulation of far-end VT Path layer PM parameters (by the near-end VT PTE) depends on whether the near-end NE supports one-bit or enhanced RDI-V, and also on which RDI-V is supported by the far-end NE.
- See page 6-123 for definitions of y, N, n, and 0.

6.2.2.9 Intermediate-Path PM

Intermediate-path PM is defined to be the transparent monitoring of a constituent channel of an incoming transmission signal by an NE that, in most cases,¹⁶ does not terminate that channel. For SONET NEs, intermediate-path PM is defined for the STS and VT Path layers, and for DSn paths. A SONET NE performing non-PJ-related intermediate-path PM examines the overhead in the monitored path and derives all of the near-end and far-end path PM parameters in each direction of transmission, while allowing the path signal to pass bidirectionally through the NE completely unaltered. The NE either frames on the incoming targeted path or uses the appropriate pointers to locate the appropriate overhead. This type of intermediate-path PM is useful (for example) when a monitored path is never terminated within a particular network provider's equipment. This is particularly true for high-capacity line services that promise certain levels of performance in terms of PM parameters such as ESs, SESSs, and UASs. In many cases, non-PJ-related intermediate-path PM is expected to be performed where SONET and asynchronous networks meet.

Figure 6-22 illustrates how non-PJ-related intermediate-path PM parameters are derived. In the figure, a bidirectional path is cross-connected through the "SONET NE", and has path terminations at PTE 1 and PTE 2. As shown in the figure, four sets of PM parameters are accumulated when intermediate-path PM is activated. For the path signal received on facility 1, there are the near-end parameters monitored on the path from PTE 1 to the SONET NE and the far-end parameters monitored via the Path REI and RDI signals carried on the path from PTE 1 to PTE 2. For the path signal received on facility 2, there are the near-end parameters monitored on the path from PTE 2 to the SONET NE and the far-end parameters monitored via the Path REI and RDI signals carried on the path from PTE 2 to PTE 1.

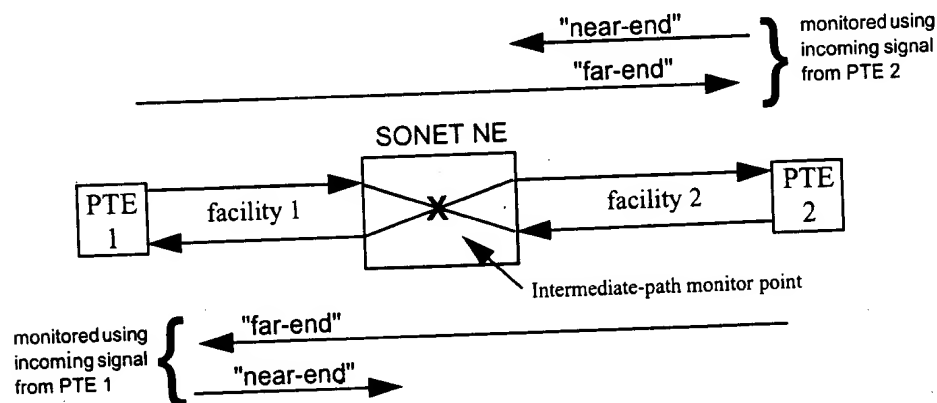


Figure 6-22. Intermediate-Path PM for Non-PJ-Related Parameters

16. An exception would be an NE that supports the accumulation of PJ-related intermediate path PM parameters at LTE for an STS Path that is subsequently terminated at STS PTE within the NE, or at STS PTE for a VT Path that is subsequently terminated at VT PTE within the NE.

Unlike the non-PJ-related parameters discussed above, the accumulation of the PJ-related parameters is affected by both the contents of the incoming signal and the actions of the LTE or STS PTE that processes the monitored path's pointers (i.e., that generates increments or decrements to frequency justify the SPE to a local clock). Therefore those parameters are typically accumulated only for the incoming path signal at the LTE or STS PTE that performs the pointer processing. This is illustrated in Figure 6-23 for two types of STS Paths (i.e., nonterminated and terminated). Note that in Figure 6-23 there is only one incoming SONET signal shown at each of the NEs' SONET "interfaces", but that additional incoming signals may be present if the NE supports line APS (e.g., linear APS or the 4-fiber BLSR architecture). In those cases it is generally expected that an NE will provide the capability to accumulate PJ-related intermediate-path PM for any of the STS Paths that make up each of the working (and possibly extra-traffic) channels, independent of which lines those channels are being selected from.

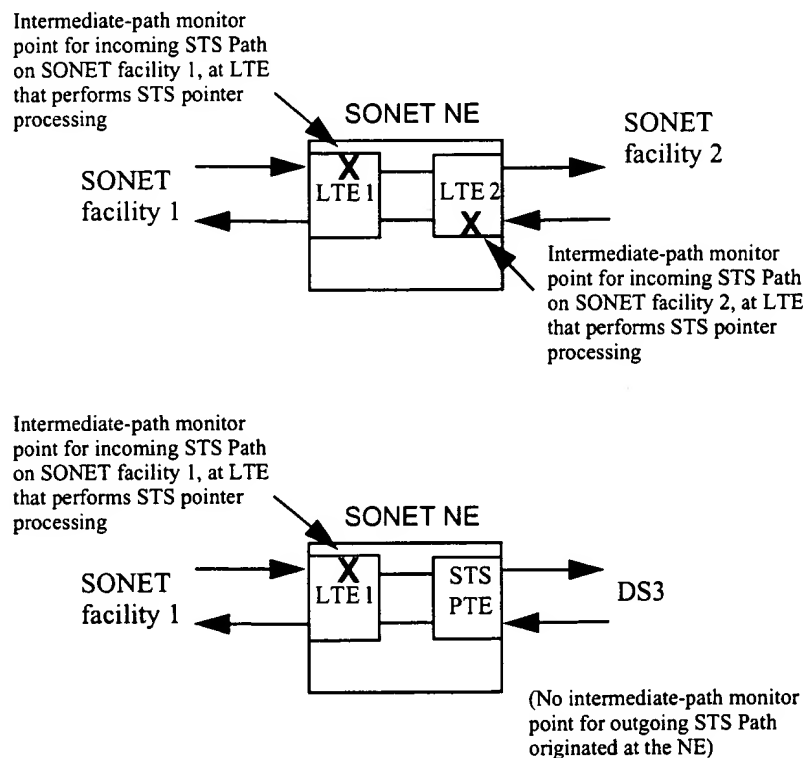


Figure 6-23. Examples of Intermediate-Path PM for STS PJ-Related Parameters

Two levels of intermediate-path PM are defined. Level-1 intermediate-path PM occurs on an incoming target path when the NE does not need to perform any additional demultiplexing of the incoming signal beyond that needed to pass the signal through the NE. Level-2 intermediate-path PM occurs on an incoming target path when a "container" path must be transparently terminated/demultiplexed on the incoming signal to monitor the desired target path.

Consistent with the description of level-1 and level-2 intermediate-path PM, two types of paths may exist at a particular NE. A level-1 path is defined to be a path that is directly visible at the NE, either at its external interfaces, or in its crossconnect fabric. A level-2 path is defined to be a path carried within a level-1 path. At a given NE, there might be several different types of level-1 and level-2 paths, depending on the types of interfaces that it has, and the types of cross-connects that it supports.

The combinations of path types and intermediate-path PM accumulation capabilities that might apply at a given SONET NE are:

- Level-1 Intermediate DS1 Path PM can occur when a DS1 is asynchronously mapped into an (originated) VT1.5 Path, and when a DS1 is asynchronously demultiplexed from a (terminated) VT1.5 Path.
- Level-2 Intermediate DS1 Path PM can occur when a VT1.5 Path carried on an incoming SONET signal and containing an asynchronously mapped DS1 is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate DS3 Path PM can occur when a DS3 is asynchronously mapped into an (originated) STS-1 Path, and when a DS3 is asynchronously demultiplexed from a (terminated) STS-1 Path.
- Level-2 Intermediate DS3 Path PM can occur when an STS-1 Path carried on an incoming SONET signal and containing an asynchronously mapped DS3 is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate VT Path PM (Non-PJ-Related Parameters) can occur when a VT Path carried on an incoming SONET signal is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate VT Path PM (PJ-Related Parameters) can occur when STS PTE processes VT pointers to frequency justify (to a local clock) a VT Path carried on an incoming SONET signal and subsequently cross-connected to an outgoing SONET signal or terminated at VT PTE within the same NE.
- Level-2 Intermediate VT Path PM (Non-PJ-Related Parameters) can occur when a VT-structured STS-1 Path carried on an incoming SONET signal is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate STS Path PM (Non-PJ-Related Parameters) can occur when an STS Path (e.g., an STS-1 or STS-3c) carried on an incoming SONET signal is cross-connected to an outgoing SONET signal.

- Level-1 Intermediate STS Path PM (PJ-Related Parameters) can occur when LTE processes STS pointers to frequency justify (to a local clock) an STS Path carried on an incoming SONET signal and subsequently cross-connected to an outgoing SONET signal or terminated at STS PTE within the same NE.

CR6-350 [685v2] A SONET NE may be required to support non-PJ-related level-1 intermediate-path PM for some or all types of level-1 paths in the NE.

CR6-351 [1091] A SONET NE containing STS PTE that processes VT pointers to frequency justify (to a local clock) VT Paths carried on an incoming SONET signal may be required to support the capability to accumulate VT PJ-related level-1 intermediate path PM parameters for those paths.

O6-352 [1092] A SONET NE containing LTE that processes STS pointers to frequency justify (to a local clock) STS Paths carried on an incoming SONET signal should support the capability to accumulate STS PJ-related level-1 intermediate path PM parameters for those paths.

CR6-353 [686] A SONET NE may be required to support level-2 intermediate-path PM for some or all level-2 path types in the NE.

R6-354 [687v2] A SONET NE that provides a given type of intermediate-path PM for a particular type of level-1 path shall provide the capability to monitor any of the level-1 paths of that type passing through the NE.

R6-355 [688v2] A SONET NE that provides intermediate-path PM for a particular type of level-2 path shall provide the capability to monitor any of the level-2 paths of that type passing through the NE.

In general, any criteria specifying that an NE must be capable of simultaneously monitoring a certain percentage or number of a particular type of path will appear in the appropriate NE-specific. However, the following requirement applies to all NEs that support intermediate path PM.

R6-356 [1093] If a SONET NE supports level-1 or level-2 intermediate path PM for a particular type of path, then the number or percentage of the paths of that type that can be simultaneously monitored shall be clearly documented.

R6-357 [689v2] A SONET NE that provides a given type of intermediate-path PM for a given path type shall provide the user the ability to activate the feature on a per-path basis. The default shall be "not active" for all paths.

- O6-358** [1094] A SONET NE that provides the capability to accumulate both PJ-related and non-PJ-related intermediate-path PM parameters for an STS or VT Path should allow the user to independently activate and deactivate the accumulation of those two sets of parameters.
- R6-359** [690] A SONET NE that provides intermediate-path PM for a given path shall not alter any bits, overhead or payload, of the monitored entity in either direction of transmission.

Note that **R6-359** [690] means an NE that provides level-2 intermediate-path PM would need to perform a non-intrusive "termination/demultiplexing" of the nonterminated container path.

- R6-360** [691v4] If a non-PJ-related intermediate-path PM feature is active for a particular path, the NE shall perform both near-end and (if defined) far-end path PM for the path signals in both directions. Except as noted below, the monitoring shall be performed as if the NE were terminating the path in each direction, using the criteria in Sections 6.2.2.5, 6.2.2.6, and 6.2.2.7 (for STS, VT, and DS_n paths).
- R6-361** [1095] If a PJ-related intermediate-path PM feature is active for a particular path, the NE shall monitor that path in the direction specified. Except as noted below, the monitoring shall be performed as if the NE were terminating the path, using the criteria in Sections 6.2.2.5 and 6.2.2.6 (for STS and VT paths).

In the preceding requirements, "monitoring" includes both the accumulation of PM parameters, and the generation of TCAs for those parameters. Note that for an NE to perform DS3 intermediate-path PM "as if it were terminating the path," the NE would need to be capable of being provisioned to expect M23 or C-bit parity application signals. That is not a capability that is typically provided in SONET NEs that transport DS3s, and therefore it is considered sufficient for the NE accumulate DS3 intermediate-path PM according to the criteria for M23 applications (for which only near-end PM parameters are defined, and for which the bit error information is provided by the P-bits). The NE may (but is not required to) allow the user to provision it to expect a C-bit parity application DS3 signal. In that case, the NE would be required to accumulate both the near-end and far-end DS3 intermediate-path PM data according to the criteria for C-bit parity applications (assuming that it is a bidirectional path).

Also note that although TIM defects may affect or contribute to the accumulation of various PM parameters at STS or VT PTE (e.g., see the parameter definitions in Sections 6.2.2.5.1 and 6.2.2.6.1), they cannot be detected at non-PTE that performs intermediate STS or VT path PM, or level-2 intermediate DS_n path PM. Therefore, the detection of TIM defects is an allowed exception to the "as if the NE were terminating the path" statements in **R6-360** [691v4] and **R6-361** [1095], and the possible impacts (both positive and negative) need to

be carefully considered before the detection of TIM defects for a path is activated. For example, if the detection of TIM-P defects is activated for a particular STS path, those defects could cause differences between the set of near-end PM data accumulated at an intermediate NE performing STS intermediate-path PM, and the set of near-end PM data accumulated at the downstream NE where the path is terminated (or the sets of far-end PM data accumulated at any NEs monitoring the upstream path).

- O6-362** [693] An NE that is capable of performing intermediate-path PM for a particular type of level-2 path should allow the user to provision whether it also performs level-1 intermediate-path PM on the container paths.
- R6-363** [694v3] An NE that is performing intermediate-path PM for a particular path shall (non-intrusively) detect and terminate AIS, LOP and RDI defects (and for DS_n paths, DS_n OOF) as if it were terminating that path.
- R6-364** [1034v2] An NE that is performing intermediate-path PM for a particular SONET path, and that supports the detection of ERDI defects for that path, shall (non-intrusively) detect and terminate UNEQ defects as if it were terminating the path.

Note that applicability of **R6-363** [694v3] and **R6-364** [1034v2] to an NE that is provisioned to accumulate only the PJ-related parameters for a particular path is currently under study (see page 6-122). In addition, there are currently no criteria in this document to indicate how an NE that supports ERDI (e.g., for paths that it terminates) and that is performing intermediate path PM on a particular path should react if it detects that the PTE at either end of that path is using one-bit RDI (instead of ERDI). This issue is discussed further in GR-253-ILR, Issue ID 253-127.

Also note that as the various PM parameters are currently defined, the detection of a PLM-P defect does not affect the accumulation of any of the STS Path layer PM parameters. Similarly, the detection of a PLM-V defect does not affect the accumulation of any of the VT Path layer PM parameters. Therefore, it is not necessary for an NE that is performing intermediate-path PM on an STS or VT path to detect and terminate PLM defects on that path. However, PLM defects detected on a container path do affect the accumulation of the target path PM parameters. Therefore, those defects are included in the following requirement.

- R6-365** [996] A SONET NE that is performing level-2 intermediate-path PM for a particular path shall (non-intrusively) detect and terminate AIS, LOP, UNEQ, PLM, and RDI defects for the container path as if it were terminating that path.

As discussed in Section 6.2.1.2.4, when an AIS, LOP, UNEQ or PLM defect is detected on a terminated STS-1 or VT path carrying an asynchronously mapped DS_n, the SONET NE is required to generate DS_n AIS downstream. In addition, the detection of DS_n AIS defects

affect the accumulation of DS_n path PM parameters (see GR-820-CORE). Therefore, when level-2 intermediate-path PM is being performed on a DS_n path, the detection of an AIS, LOP, UNEQ or PLM defect on the container path must affect the accumulation of the PM parameters in the same way that the detection of DS_n AIS would.¹⁷

R6-363 [694v3] and R6-365 [996] use the phrase "an NE that is performing" which means they are applicable when an intermediate-path PM feature is active for the particular path. Note that it is also acceptable for an NE that is capable of performing intermediate-path PM on a particular path, but that is currently not provisioned to do so, to (non-intrusively) detect and terminate defects as if the intermediate-path PM feature were active.

Finally, see Section 6.2.1.8.7 for criteria concerning the declaration and clearing of failures based on the defects detected and terminated according to the preceding criteria.

6.2.3 Testing Process

Testing deals with procedures that result in isolation of a failure to a replaceable or repairable entity. Maintenance tools that achieve this isolation, besides those built into the SONET signal format, are test access, diagnostics, and loopbacks.

The long-term objective for testing is to evolve toward self-diagnosing NEs. Steps toward this goal are effected in the alarm and PM plans. Another step is reflected in the requirements for diagnostics, discussed in Section 6.2.3.2. Testing activities include:

- Analyzing alarms, PM data, and maintenance signals
- Executing diagnostics
- Executing controls, such as switching to protection
- Activating loopbacks
- Test access for signal measurements.

In each of these activities, operations personnel gain access through either a local craftsperson interface or the remote operations interface.

6.2.3.1 Test Access

Test access allows access to a signal for the purposes of non-intrusive monitoring and intrusive testing. For SONET NEs, this access falls into three possible categories: access

17. Note that the detection of a TIM-P defect (if activated) or TIM-V defect (if defined and implemented) would also be expected to cause the generation of DS_n AIS downstream. However as discussed above, those defects cannot be detected at non-PTE, and therefore they cannot affect the accumulation of the DS_n path PM parameters (in the same way as DS_n AIS) when level-2 intermediate DS_n path PM is being performed.

to the fiber for monitoring and testing the optical signal and the fiber; access to the SONET signal for monitoring and testing the format, mapping, and equipment specifications; and digital test access to lower-speed digital signals to test the service.

6.2.3.1.1 *Fiber Access*

GR-1309-CORE, *TSC/RTU and OTAU Generic Requirements for Remote Optical Fiber Testing*, describes functional requirements for fiber testing.

- R6-366** [706] Access to the fiber for fiber testing (e.g., for identifying the location of a fiber break) shall be provided.

This fiber access currently implies disconnecting the optical source or receiver.

- O6-367** [707] To aid in the remote sectionalization of a problem to the fiber or to one of the terminals, the ability to switch the fiber to an alternate source or an alternate receiver should be provided.

6.2.3.1.2 *SONET Signal Test Access*

This section provides the following objective for local test access until more experience is gained with SONET signals.¹⁸

- O6-368** [709] A SONET NE should provide local test access to individual STS-1s and STS-3s within an OC-N or STS-N electrical signal.

If an NE provides local test access, the following requirements apply:

- R6-369** [710] The SONET electrical level test access shall be in accordance with the specifications defined in Section 4.4 for STS-1 and STS-3 electrical signals.
- R6-370** [711] The SONET electrical level test access shall provide a non-intrusive and hitless monitor mode.
- R6-371** [712] The SONET electrical level test access shall provide the ability to perform intrusive tests.
- R6-372** [713] When intrusive testing occurs, the NE shall provide the appropriate Path AIS in the non-test direction.

18. The support of remotely controlled access and testing for OC-N signals, and requirements for remotely controlled access and testing of SONET electrical interfaces, are for further study.

- R6-373** [714] In NEs supporting APS, the test access shall not impair the working channel on the protection line.

6.2.3.1.3 Digital Test Access

SONET NEs that are VT (or DS1) programmable (i.e., that allow the flexible assignment of low-speed signals to timeslots within an OC-N signal) may provide DS1 remote test access.

- CR6-374** [715] A SONET NE that is VT programmable may be required to provide DS1 remote test access.

This conditional requirement may be changed to a requirement in SONET NE-specific GRs, TRs, or TAs. If a SONET NE provides DS1 remote test access, the following requirements apply:

- R6-375** [716] The DS1 test access shall meet the criteria in FR-476, *OTGR Section 6: Network Maintenance: Access and Testing*, unless deviations from this are explicitly identified within NE-specific GRs, TRs, or TAs.
- R6-376** [717] When TL1 interfaces are used, SONET NEs shall use the TL1 messages of GR-834-CORE, *OTGR Section 12.4: Network Maintenance: Access and Testing Messages*, for any test access functions provided.
- R6-377** [718] When CMISE/OSI interfaces are used, SONET NEs shall adhere to the object model and use the CMISE service mappings of GR-1031-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools: Metallic Test Access*, for any test access functions provided.

6.2.3.2 Diagnostics

The term diagnostic, as used in this section, has several connotations. There are those diagnostics that the equipment manufacturer provides for internal troubleshooting. Some diagnostics in this category may run continuously, some may run on a designed schedule (e.g., once an hour or triggered by an event), and some may run on a schedule the user is permitted to specify. Typically, these routine diagnostics are design dependent, and during normal operation the user is unaware of them. The ability to manually initiate routine diagnostics between the scheduled times is referred to as "on-demand". The second sense in which the term diagnostics is used is as a user's testing tool. These tools are functions that are activated on demand (initiated at a WS or OS interface) to retrieve information from the overhead or operate special circuitry to support trouble analysis (e.g., the corrupted BIP) or to run an NE diagnostic and see the result.

- R6-378** [719] A SONET NE shall provide diagnostic capabilities. As a minimum, diagnostics shall be provided to detect the equipment failures listed in Section 6.2.1.1.4. The diagnostic capabilities shall run routinely and on demand. When these diagnostics are run on demand, the NE shall provide the user with the results.
- O6-379** [720] An NE should support additional diagnostics that provide the ability to isolate an equipment trouble to the replaceable circuit pack or module. These diagnostics should not interfere with working services and may be run routinely or on demand.
- R6-380** [721] Diagnostics that interfere with service shall not run routinely unless permitted by the user.¹⁹

6.2.3.2.1 Physical Layer

As discussed in Section 6.2.2.2.2, in some situations it may be useful for a user to be able to determine the current value of an optical transmitter's LBC or OPT, or an optical receiver's OPR. If an NE supports the Physical layer PM parameters defined in Section 6.2.2.1 and meets **O6-315 [1084]**, then this capability can be provided as a part of its Physical layer PM feature. On the other hand, if the NE does not support one or more of the Physical layer PM parameters or does not meet **O6-315 [1084]**, then the following objectives apply.

- O6-381** [722v2] An NE that does not support the OPR_{normal} PM parameter for an optical receiver should provide an on-demand diagnostic that reports the received optical power OPR (not normalized), preferably in units of dBm.
- O6-382** [723v2] An NE that does not support the OPT_{normal} PM parameter for an optical transmitter should provide an on-demand diagnostic that reports the transmitted optical power OPT (not normalized), preferably in units of dBm.
- O6-383** [724v2] An NE that does not support the LBC_{normal} PM parameter for an optical transmitter should provide an on-demand diagnostic that reports the laser bias current LBC (not a normalized), preferably in microamperes (μA).
- O6-384** [1096] An NE that supports the OPR_{normal} PM parameter but does not meet **O6-315 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of OPR_{normal} .

19. It is expected that the supplier would provide a complete list of any diagnostics that interfere with service.

- O6-385** [1097] An NE that supports the OPT_{normal} PM parameter but does not meet **O6-315** [1084] should provide an on-demand diagnostic that reports the current value (not the snapshot value) of OPT_{normal} .
- O6-386** [1098] An NE that supports the LBC_{normal} PM parameter but does not meet **O6-315** [1084] should provide an on-demand diagnostic that reports the current value (not the snapshot value) of LBC_{normal} .

6.2.3.2.2 Section Layer

The criteria that appear in this section describe a diagnostic that involves looping back a SONET signal immediately before it is transmitted, to verify the integrity of the electronics associated with the receiver or the transmitted signal (see Figure 6-24). In Issue 1 of this document, these criteria appeared in Section 6.2.3.3.1, and the feature was identified as a SONET Terminal Loopback. However, the term “loopback” is generally associated with features that involve the use of external test equipment to monitor the looped back signal. The feature described here is strictly internal to the SONET NE at which it is performed, and therefore is considered to be a diagnostic.

- R6-387** [740v2] The SONET NE shall provide the functional equivalent of the diagnostic illustrated in Figure 6-24.

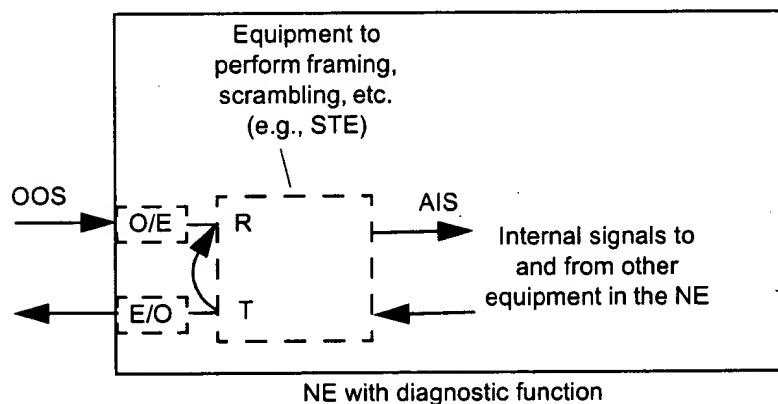


Figure 6-24. Section Layer Diagnostic

In this diagnostic, the signal is looped back by connecting the outgoing signal immediately before the electrical-to-optical conversion (after scrambling) to the associated receiver. In

the case of an electrical interface (i.e., if there is no electrical-to-optical conversion), the loop occurs after scrambling. Note that this diagnostic will interrupt an incoming SONET signal, and therefore it must not be initiated by a command carried in the Section DCC of that signal (because it will interrupt communications between the NE performing the loopback and the user). The use of a DCC in this manner could cause the diagnostic to be initiated without providing a means for it to be terminated. Also note that if the NE supports linear APS at the interface where the diagnostic is being performed, then the traffic could be selected from the line that is not affected by the diagnostic and no service-interruption need occur.

- R6-388** [741] An NE shall deny a request for this diagnostic if the incoming facility associated with the receiver is In_Service or Out_of_Service-autonomous, as defined in GR-1093-CORE.
- R6-389** [743] While the diagnostic is being performed, the NE shall recognize that it is in a test state, as defined in GR-1093-CORE, and it shall not take action to revert to a previous state.
- R6-390** [745] The NE shall exit the test state, as defined in GR-1093-CORE, when the diagnostic is deactivated.
- R6-391** [744] An NE shall notify the OS when the diagnostic has been activated and when it has been deactivated.
- R6-392** [746] The NE shall determine if the receiver is able to frame on the delivered signal.
- R6-393** [747] The NE shall detect errors in the Section BIP-8 (B1) byte, or if the B1 byte is not processed, in the Line BIP-8s (B2).
- R6-394** [748v2] The NE shall report results indicating that the equipment could not frame, that CVs were detected, or that no framing problems or CVs were detected.

6.2.3.2.3 *Signal Identification*

This section contains criteria for diagnostics that allow a signal to be traced back to its source (e.g., for troubleshooting purposes). Diagnostics are defined that use the STS Path Trace (J1 byte), and the STS and VT Signal Labels (C2 bytes and V5 bits 5 through 7).

6.2.3.2.3.A STS Path Trace

The STS Path Trace may be used for connectivity verification even in applications where the detection of TIM-P defects (see Section 6.2.1.1.9) is not used. The following criteria, along with the criteria in Section 6.2.1.1.9 related to the provisioning and transmission of path trace strings, are intended to support that use.

- R6-395** [728] STS PTE shall provide an on-demand diagnostic to detect and report the contents of the received STS Path Trace.
- CR6-396** [729] LTE with STS cross-connection capabilities may be required to provide an on-demand diagnostic to detect and report the contents of the STS Path Trace in the (nonterminated) STS path designated by the user.
- O6-397** [730] STS PTE should provide a diagnostic that, when activated, continuously monitors the incoming STS Path Trace.

In addition, LTE with STS cross-connection capabilities may (but is not required to) provide a diagnostic that, when activated, continuously monitors the incoming STS Path Trace in nonterminated STS paths.

As discussed in Section 6.2.1.1.9 for the transmitted STS Path Trace, it is assumed that the path trace strings detected by SONET NEs that provide these diagnostics will generally consist of ASCII printable, NULL, <CR>, and <LF> characters. Most of the following criteria are based on that assumption.

In general, it is possible for a SONET NE that supports on-demand STS Path Trace diagnostics to simply report the detected path trace, or to also report whether the detected path trace matches a user-provisioned "expected" path trace. Similarly, a SONET NE that supports a continuous STS Path Trace diagnostic could compare the current incoming path trace to either a previously received path trace or to a user-provisioned expected path trace.

- CR6-398** [999v2] A SONET NE may be required to support a feature to allow the user to provision, for diagnostics purposes, the "expected" path trace for each STS path that it terminates.

Note that support of a feature that allows the user to provision the "expected" path trace for TIM-P defect detection purposes is required in Section 6.2.1.1.9. Also note that the use of an expected path trace feature is not recommended for STS paths that are not terminated. The use of such a feature could result in excessive reprovisioning at intermediate NEs when the two end-points (i.e., the STS PTE) are changed.

- R6-399** [1000v2] If an NE supports a feature that allows the user to provision an expected path trace for diagnostics purposes and that feature uses ASCII characters, then the following apply:

- The feature shall allow the user to enter a string of up to 62 characters

- ...
- The feature shall place no restriction on the contents of the string, except that the characters shall be ASCII printable characters.
- R6-400** [731v3] A SONET NE that is continuously monitoring an incoming path trace and does not support an expected path trace feature for diagnostics purposes shall compare the incoming path trace with a previously received path trace.
- CR6-401** [1002v2] A SONET NE that is continuously monitoring an incoming path trace and supports an expected path trace feature for diagnostics purposes, but that is not provisioned with an expected path trace, may be required to compare the incoming path trace with a previously received path trace.
- R6-402** [1003] As a minimum, an NE that compares the incoming path trace with a previously received path trace shall be capable of performing a case-sensitive comparison of ASCII-based path traces.
- O6-403** [1004] An NE that compares the incoming path trace with a previously received path trace should be capable of performing that comparison on a byte-by-byte basis (i.e., independent of whether the contents are ASCII-based or contain <CR> and <LF> characters).

For a SONET NE that meets **O6-403** [1004], neither the incoming path trace or the previously received path trace would need to consist of ASCII printable characters. Therefore, such an NE would be capable of receiving an SDH 16-byte E.164 path trace (repeated four times) and detecting if that path trace changes.

- R6-404** [735v3] A SONET NE that is monitoring an incoming path trace for changes or mismatches for diagnostics purposes shall suspend the STS Path Trace monitoring function when the J1 byte cannot be accessed (e.g., when an LOP-P or AIS-P defect has been detected by the STS PTE). If the NE is monitoring for changes in the incoming path trace, the contents of the path trace before the disruption shall be used as the starting value following a restart.
- R6-405** [736v2] A SONET NE that has suspended monitoring of an incoming path trace because the J1 byte could not be accessed shall resume monitoring when the J1 byte can again be accessed.
- R6-406** [732v2] A SONET NE that is monitoring for changes of the incoming path trace shall detect when a sustained change in the path trace content occurs. Upon detecting a sustained change, the NE shall send a message to an OS. The level of the message shall be Not Alarmed, and it shall

include both the previously received path trace, and the new path trace (assuming they are ASCII-based).

R6-407 [1005v2] A SONET NE that is monitoring for a mismatch between the incoming path trace and an expected path trace for diagnostics purposes shall detect when a sustained mismatch occurs. Upon detecting a sustained mismatch, the NE shall set an indication for that path and send a message to an OS. The default level of the message shall be Not Alarmed, and it shall include both the expected path trace and the new path trace (assuming they are ASCII-based).

R6-408 [1006v2] A SONET NE that is monitoring the incoming path trace for diagnostics purposes and that has detected a sustained mismatch shall detect when the incoming path trace matches the expected path trace. Upon detecting a match, the NE shall clear the indication for that path and send a clear message to the OS (if the mismatch was reported to an OS).

A sustained change or mismatch of the STS Path Trace is one where the new path trace is consistently being received, as opposed to a short-term disruption or mismatch due to (for example) a burst of errors. Note that a change in the "starting position" of the incoming path trace is not considered a sustained change or mismatch. Such a change could be caused by (for example) an upstream protection switch where transmission delay differences cause a J1 byte to appear to be dropped or repeated in successive frames (so that one path trace string would appear to contain 63 or 65 bytes). Also note that **R6-406 [732v2]**, **R6-407 [1005v2]**, and **R6-408 [1006v2]** do not specify a time period within which an NE must report that it has detected a change or mismatch. In general, the goal is for such events to be reported without excessive delays, but for the detection mechanism to be robust in the presence of performance degradations. Timing similar to that required for the defects and failures defined in Section 6.2.1 (e.g., report a change or mismatch if it persists for approximately 2.5 seconds) would be appropriate.

CR6-409 [1007v2] An NE that monitors the incoming path trace for mismatches for diagnostics purposes may be required to be user-provisionable to report a detected mismatch as alarmed or Not Alarmed.

Note that if an NE is monitoring a particular path for changes rather than mismatches, no condition exists that would cause it to send a clear message after it has reported an STS Path Trace change (e.g., a detected change back to the old path trace would cause another change to be reported, not a clear message). Therefore, an STS Path Trace change cannot be alarmed. In addition, since a continuous STS path trace monitoring diagnostic is a feature that must be activated by the user, it is assumed that the user will want any detected changes or mismatches to be autonomously reported. Therefore, it is not necessary for an NE to allow the user to provision a change or mismatch to be Not Reported.

6.2.3.2.3.B STS and VT Path Signal Label

STS and VT Signal Label diagnostics can be used to identify the construction of STS and VT SPEs for trouble-shooting purposes.

- R6-410** [737] A SONET NE shall provide an on-demand diagnostic to report the value currently being received in the STS and VT Signal Labels at STS PTE and VT PTE, respectively.
- R6-411** [1008] A SONET NE that supports the detection of PDI-P defects (e.g., at an STS-level path selector in a UPSR NE) shall provide an on-demand diagnostic to report the value currently being received in the STS Signal Label contained in each STS path that it is monitoring.
- CR6-412** [1009] An NE that supports the generation of PDI-P signals may be required to support an on-demand diagnostic that reports the code that is currently being inserted on the specified STS Signal Label.

6.2.3.2.4 Error Monitoring

- R6-413** [738] A SONET NE shall provide the ability to transmit, on demand, a fully corrupted BIP value (all parity check bits inverted) for the Section, Line, STS Path, or VT Path (as appropriate to the NE). The NE shall provide the user the capability to specify the approximate duration of the corrupted BIP diagnostic as some number of units of time or some number of frames (or VT superframes for VT Path BIP).

The above demand diagnostic can be used (for example) to verify that the performance monitoring functions of Section 6.2.2 in another NE are working properly.

6.2.3.3 Loopbacks

To support pre-service operations practices and test-related activities in some applications, SONET NEs may need to provide loopbacks for SONET and DS_n signals. Two types of SONET and DS_n signal loopbacks are discussed in this document. These are terminal loopbacks, and facility loopbacks. In general, loopbacks interrupt the flow of traffic, change the normal transmission, and often require coordinated activity as two or more NEs are affected. Because of this potential impact on the network, the use of a loopbacks in the SONET network as routine practice is discouraged. In addition, a DCC must not be used to initiate a SONET loopback if that loopback will interrupt communications between the NE performing the loopback and the user. The use of a DCC in this manner could cause a loopback to be initiated without providing a means for it to be terminated.

6.2.3.3.1 SONET Terminal Loopback

In general, loopbacks involve the use of external test equipment to monitor the looped back signal. Conversely, the feature that was described in this section in Issue 1 of this document was strictly internal to the SONET NE. Therefore, the criteria related to that feature have been moved to Section 6.2.3.2.2 and there currently are no applicable "SONET terminal loopback" criteria. Although no such criteria are expected to be added to this document in the future, several issues related to terminal loopbacks are discussed below.

In a terminal loopback, the signal that is about to be transmitted is connected to the associated, incoming receiver. For example, a SONET regenerator could provide a terminal loopback capability such that an incoming signal is looped back at the "far side" of the regenerator (see Figure 6-25). Such a capability could allow a user to test a SONET line system step-by-step (i.e., test the facilities between the test equipment and regenerator "A" using the SONET facility loopback described in Section 6.2.3.3.2 at "A", test those facilities and "A" using the terminal loopback capability at "A", test the additional facilities between "A" and regenerator "B" using the SONET facility loopback at "B", etc.). Note that if an NE terminates the SONET Line layer, it would generally not be recommended that it perform terminal loopbacks at its high-speed SONET interfaces. Such loopbacks could be used to determine if the NE is multiplexing and demultiplexing a particular low-speed signal (e.g., a DS_n) correctly; however, they would also disrupt all of the other low-speed signals carried on the looped back signal. Conversely, it may be appropriate for an NE that terminates the SONET Line layer to provide terminal loopback capabilities for its low-speed SONET interfaces.

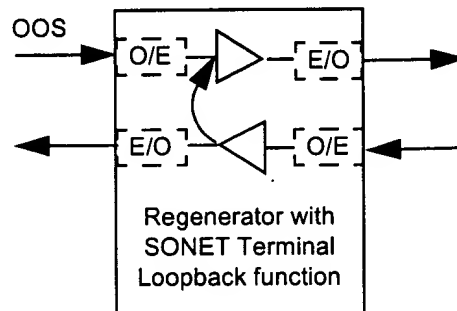


Figure 6-25. SONET Terminal Loopback

6.2.3.3.2 SONET Facility Loopback

In general, a facility loopback connects the incoming received signal to the transmitter in the return direction (see Figure 6-26).

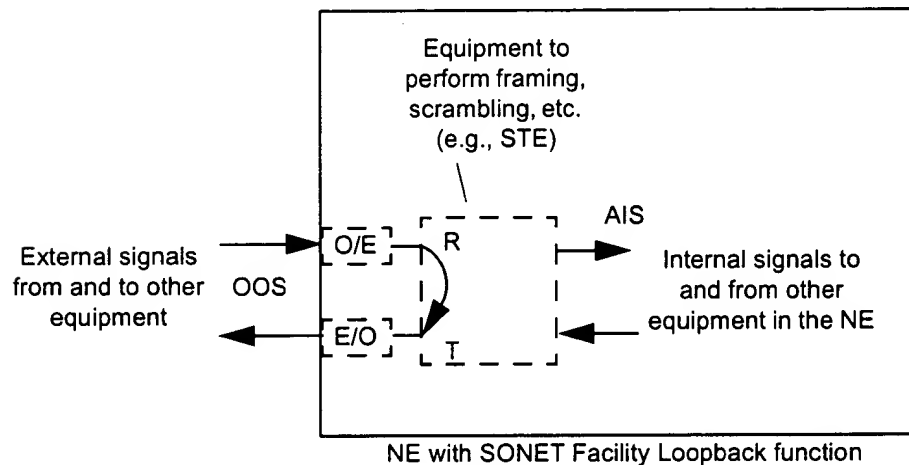


Figure 6-26. SONET Facility Loopback

O6-414 [749] A SONET NE should provide a SONET facility loopback, as illustrated in Figure 6-26.

In a SONET facility loopback, the signal is looped back by connecting the incoming received signal immediately following the optical-to-electrical conversion (before descrambling) to the associated return transmitter. In the case of an electrical interface (i.e., if there is no optical-to-electrical conversion), the loop occurs before descrambling.

R6-415 [750] A SONET NE shall deny a SONET facility loopback request if either the impacted direction of transmission (i.e., the direction of the looped signal) or the incoming signal to be looped is In_Service or Out_of_Service-autonomous, as defined in GR-1093-CORE.

R6-416 [751v2] When a SONET facility loopback is activated, the SONET NE shall place the associated out-of-service facilities into a test state, as defined in GR-1093-CORE.

- R6-417** [756] The SONET NE shall exit the test state, as defined in GR-1093-CORE, when the loopback is deactivated.
- R6-418** [755] The SONET NE shall notify the OS when the loopback has been activated and when it has been deactivated.

6.2.3.3.3 *DSn Loopback*

- O6-419** [757v2] A SONET NE providing DS_n line terminations should provide DS_n terminal loopback capabilities, as shown in Figure 6-27.
- O6-420** [1010] SONET NE providing DS_n line terminations should provide DS_n facility loopback capabilities, as shown in Figure 6-28.

Note that in the DS_n terminal loopback, the DS_n is looped back toward the SONET system just before being transmitted on a DS_n line. This DS_n can be monitored (for example) at the point where the DS_n entered the SONET system as a check on the performance of that system. In the DS_n facility loopback, on the other hand, the DS_n is loopback immediately after entering the SONET system. Therefore, it can be used as a check of the performance of the DS_n facility. Both of these DS_n loopbacks are initiated and removed by commands sent to the NE (e.g., from a craft interface or an OS, either directly or via an NE/NE communications channel such as the Section DCC).

In addition to the types of DS_n loopbacks covered by the preceding criteria, an NE may support DS_n loopbacks that are activated by messages carried in the DS_n signals [e.g., by ESF Data Link messages in an ESF DS1, or by Far End Alarm and Control (FEAC) signals in a C-bit parity DS3]. See GR-499-CORE for additional information concerning those types of DS_n loopbacks. In addition, SONET NE-specific GRs may contain additional criteria on DS_n loopback capabilities.

NE with DS_n Terminal Loopback function

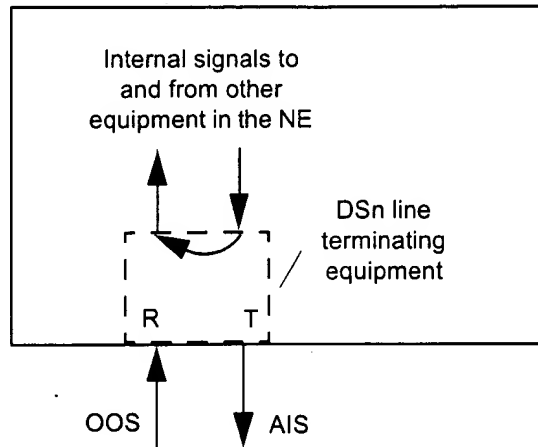


Figure 6-27. DS_n Terminal Loopback

NE with DS_n Facility Loopback function

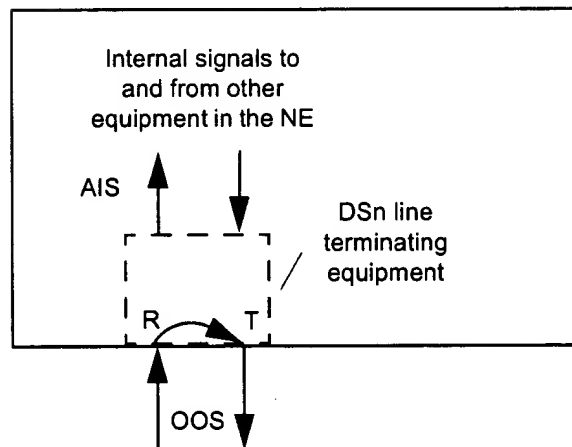


Figure 6-28. DS_n Facility Loopback

6.2.4 Control Features

This section highlights control features required for SONET NEs that are necessary to maintain the NE. Other control features explicitly mentioned earlier in Section 6 are not restated.

R6-421 [759] The following control functions shall be provided:

- ... 1. Reinitialize system – System reinitialization, or “hard boot,” reloads the operating system of the NE and may affect the state of memory and other resources.
- ... 2. Restart system – System restart, or “soft boot,” reloads only an application onto the system and not the operating system. System restart should not affect the state of nonvolatile memory and other resources. Failure states, protection switching configuration, performance parameters, and other information necessary for fault isolation should not be affected.
- ... 3. Reestablish a failed entity – Reestablishing failed entity involves temporary techniques to restore a service when an entity has failed. For example, when a facility fails, restoration of a service may involve reconfiguring routing tables or rerouting facilities.
- ... 4. Remove an entity from service to run tests – For such tests, traffic should be switched off the entity, and subsequent indications should be suppressed.
- ... 5. Inhibit and allow alarmed and nonalarmed indications – This capability permits the user to suppress and restart messages from the NE.
- ... 6. Check status of equipment configuration – Equipment configuration status shall be retrievable. Equipment configuration status includes the indication of the active hardware entities in replicated equipment and the active synchronization source.
- ... 7. Protection switch capabilities – See the required functions identified in Section 5.3.6.
- ... 8. Manual switch from active to standby – This capability is used for any replicated hardware or software.
- ... 9. Manual switch between synchronization sources – See the required functions identified in Section 5.4.6.

R6-422 [760] A SONET NE shall notify the OS when any control function is executed.

7. Other Generic Criteria

7.1 Physical and Environmental Criteria

This section provides references for physical and environmental criteria for equipment and for outside plant cable.

7.1.1 Operational Environment for Equipment

The operational environment is a set of conditions of temperature, humidity, and airborne contaminant concentration over which the specified parameters maintain their stated performance rating. The requirements differ, depending on the type of structure that houses the equipment. GR-63-CORE, *Network Equipment - Building System (NEBS) Equipment Requirements: Physical Protection*, provides physical and environmental requirements for central office equipment and remote terminal equipment installed in environmentally controlled structures such as Controlled Environment Vaults (CEVs). TA-NWT-000487, *Generic Requirements for Electronic Equipment Cabinets*, contains the technical requirements for structures with limited or no environmental controls. TA-NPL-000286, *NEBS Generic Engineering Requirements for Systems Assembly and Cable Distribution*, contains further information on engineering and installation of central office equipment.

For equipment installed in environmentally controlled structures such as central offices or CEVs, the temperature and humidity conditions are presented in GR-63-CORE. Temperature and humidity ranges are given for both normal operating conditions and for short-term conditions when temperature and humidity limits are more severe than normal. GR-63-CORE also presents conditions for the central office levels of airborne contamination, which also apply to CEVs.

TA-NWT-000487 details the environmental stresses under which the equipment installed within structures with limited or no environmental controls is expected to function. For equipment housed in structures with limited environmental controls, the temperature and humidity limits are given in Section 4.1.3 of TA-NWT-000487. Sections 9 and 10 of TR-NWT-000057, *Functional Criteria for Digital Loop Carrier Systems*, provide other physical and environmental requirements for Digital Loop Carrier (DLC) equipment. Many of the requirements presented in those two sections refer to GR-63-CORE as the source for environmental requirements for network equipment. For example, Section 4.6 of GR-63-CORE presents the airborne contaminant requirements for outdoor, urban levels. These three documents are used together to properly determine the compliance of structures with limited environmental controls and the equipment housed therein.

7.1.2 Outside Plant Cable

GR-20-CORE contains criteria for physical and environmental requirements for outside plant cable.

7.2 Equipment Design

- O7-1** [761] The equipment should be designed to minimize the investment in the frame and bay-work by using the modular design concept to minimize the costs associated with installation and the ongoing operation of the system. TR-NWT-000078, *Generic Physical Design Requirements for Telecommunications Products and Equipment*, contains physical design requirements.

7.3 Documentation and Training

The BCC may assume responsibility for engineering, constructing, and installing the transmission system.

- R7-2** [762] The supplier shall provide appropriate documentation and training. To accomplish this in a safe and cost-effective manner, criteria for supplier documentation for NEs in TR-TSY-000454, *Supplier Documentation for Network Elements*, shall be followed. Criteria for supplier documentation for outside plant cable are in GR-20-CORE.
- R7-3** [763] The supplier shall provide training in accordance with TR-OPT-000839, *Supplier-Provided Training Generic Requirements*.

TR-OPT-000839 provides requirements for:

- Existing training
- Training to be developed by suppliers
- Course content and training documentation
- Course delivery
- Training maintenance and updating
- Product support.

7.4 Safety

This sections refers to station equipment safety and fiber optic cable safety.

7.4.1 Station Equipment Safety

In general, most of the safety-related criteria applicable to SONET NEs can be found in Sections 12 and 14 of GR-499-CORE, which covers areas such as safety labels (format and locations), and user access to high voltages or temperatures. In addition to those criteria, the requirements in this section are also applicable to SONET NEs.

- R7-4** [764] All safety labels shall be visible to craftspersons when equipment covers are in place and when they are removed.

To minimize the exposure of personnel to hazardous voltages, the following requirements are applicable:

- R7-5** [765] Voltages at or above 140 V_{dc} or 50 V_{rms ac} shall be enclosed or guarded to prevent contact. Safety labels shall be conspicuous when the guards are in place and when they are removed.
- R7-6** [766] The design shall allow craftspersons safe access to parts if metal tools are to be used (e.g., insulating sleeves to guide screwdrivers to recessed potentiometers when nearby parts have hazardous voltages present).
- R7-7** [767] Arrangements shall be provided to discharge large capacitors (e.g., "bleeder" resistors).
- R7-8** [768] All external metal parts shall be grounded.

For additional information on grounding, refer to GR-63-CORE.

- R7-9** [771] The fiber optic system and required optical test equipment shall be registered and certified with the Department of Health, Education and Welfare Bureau of Radiological Health as specified in 21 CFR 1040.10, *Performance Standard for Laser Products*. Documentation demonstrating system certification shall be available to assist in the determination of fiber optic safety precautions required to install, operate, and maintain the system.
- R7-10** [772] The equipment involved shall conform to the applicable performance requirements, labeling requirements, and informational requirements as specified in 21 CFR 1040.10.

7.4.2 Fiber Optic Cable Safety

- R7-11** [773] The fiber optic cable and required optical splicing and test equipment shall be registered and certified with the Department of Health, Education and Welfare Bureau of Radiological Health as specified in 21 CFR 1040.10. Documentation demonstrating system certification shall be available to assist in the determination of fiber optic safety precautions required to install, operate, and maintain a fiber optic system.
- R7-12** [774] The equipment involved shall conform to the applicable performance requirements, labeling requirements, and informational requirements as specified in 21 CFR 1040.10.

7.5 Quality and Reliability

This section refers to the quality and reliability of station equipment and fiber optic cable.

7.5.1 Network Equipment Reliability

Reliability requirements are intended to help ensure that a product will be able to meet all technical specifications consistently and throughout the lifetime of the product. These criteria cover the areas of system availability, system qualification, manufacturing tests and inspections, software reliability, and customer support and field reliability. TR-NWT-000418 details these requirements.

7.5.2 Fiber Optic Cable Quality and Reliability

General quality and reliability requirements for fiber optic cable are given in GR-20-CORE. They deal with documentation, manufacturing program analysis, quality surveillance, process and product verification, and initial product qualification and periodic requalification.

7.5.3 Component Reliability Assurance

Component reliability assurance criteria address the necessary attributes and minimum practices of an equipment supplier's component qualification and lot control procedures. TR-NWT-000357, *Generic Requirements for Assuring the Reliability of Components Used in Telecommunication Systems*, details criteria for general components. TR-NWT-000468, *Reliability Assurance Practices for Optoelectronic Devices in Central Office Applications* and TA-NWT-000983, *Reliability Assurance Practices for*

Optoelectronic Devices in Loop Applications, address optoelectronic components in central office and loop applications, respectively. Hybrid microcircuits are the subject of TR-NWT-000930, *Generic Requirements for Hybrid Microcircuits Used in Telecommunications Equipment*.

7.6 Other Requirements and Objectives

7.6.1 Human Factors

Human factors requirements and objectives are specified in Section 12 of GR-499-CORE. In addition, FR-480, *OTGR Section 10: User System Interface*, contains criteria concerning the craft interface.

7.6.2 Technical Analysis

- O7-13 [775] It is an objective that suppliers allow Bellcore to analyze products and processes to determine their products' compliance with this document.

8. SONET Operations Communications

SONET NE operations communications criteria are consistent with the Telecommunications Management Network (TMN) concept in ITU-T M.3010 (1996) *Principles for a telecommunications management network*. A TMN is a support network that provides operations communications paths for SONET Operations System/Network Element (OS/NE), Mediation Device (MD)/NE, NE/NE, and Work Station (WS)/NE communications. This section briefly describes the role of SONET NEs in a TMN, and focuses on the implementation of the communications network functions and mediation functions by using SONET NEs and SONET overhead channels, specifically SONET Data Communications Channels (DCCs). The use of other technologies including Local Area Networks (LANs) and Mediation Devices (MDs) for mediation functions is also briefly discussed.

Operations communications criteria for SONET NEs depend on the location of the NE in the TMN architecture. More specifically, operations communications criteria depend on whether the NE serves as a Gateway NE, Intermediate NE, or End NE, as Figure 8-1 shows.

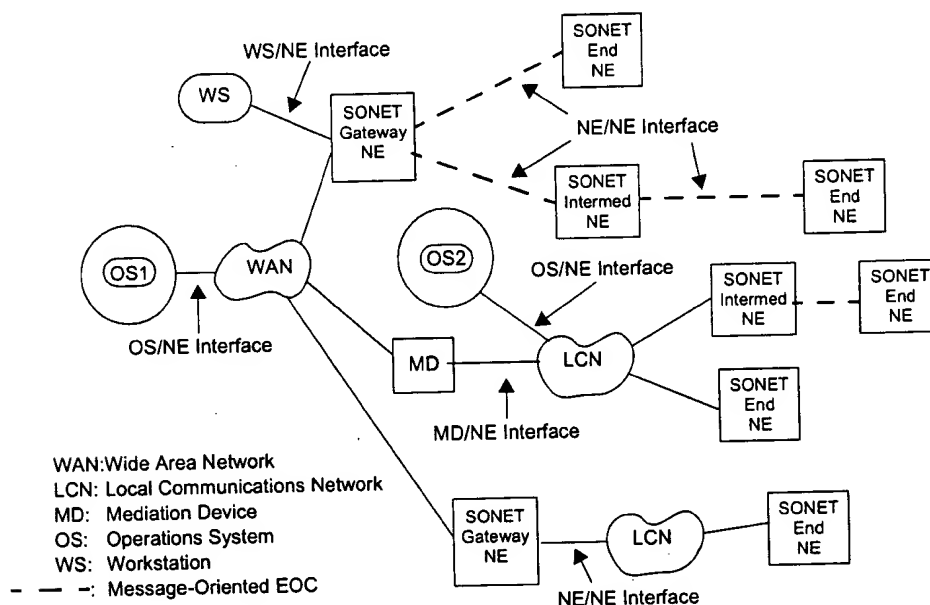


Figure 8-1. SONET Operations Communications: Example NE and Interface Types

Section 8.1 discusses the SONET Operations Communications Architecture, including Gateway NEs, Intermediate NEs, End NEs, MDs, and the operations communications specifications associated with them. Section 8.2 describes four types of communications

supported (OS/NE, MD/NE, NE/NE, and WS/NE communications). These types of communications are provided by using any combination of OS/NE, MD/NE, NE/NE, and WS/NE interfaces. Section 8.3 provides requirements for the operations communications interface (OS/NE and NE/NE). Section 8.4 discusses interworking between OSs and SONET NEs. Section 8.5 discusses SONET Operations Communications Routing. Section 8.6 provides requirements for Craftsperson/NE interfaces, and Section 8.7 provides requirements for the TL1 TID Address Resolution Protocol (TARP).

8.1 SONET Operations Communications Architecture

8.1.1 Architecture Overview

SONET operations communications architectures will vary depending on configuration (e.g., communications within a site or between sites) and application (e.g., OS-NE or NE-NE, IEC-LEC, survivable rings). This section will look at some operations communications architectures that may be used by network providers. Within a site, typically drop-side SONET interfaces will be used between connected SONET NEs; thus, no DCC will be supported. In this case, a Local Area Network (IEEE 802.3 LAN) can provide an alternate means for intra-site operations communications. However, there may be cases where a line-side interface (i.e., with the DCC) is required by a network provider for an intra-site transport connection. One such example may be for transport connections between exchange carriers where operations communications security is a concern. Another example of where a DCC may be used for intra-site communications is a Controlled Environment Vault (CEV) where only a few SONET NEs may reside. Therefore, for certain applications, security, reliability, or cost may make the use of the DCC rather than an intra-site LAN a better choice for intra-site operations communications.

Figure 8-2 shows a generalized view of a SONET operations communications architecture that includes a Wide Area Network (WAN), a LAN, and DCC tree and ring connections. Figure 8-3 shows the protocol stacks for the X.25 WAN, the DCC, and the LAN. In all three cases, the Connectionless Network layer Protocol (CLNP ISO 8473) resides at layer 3. Thus, interworking the three protocol stacks is done by standard routing and relaying functions. (Section 8.5 contains specific requirements for support of routing protocols.

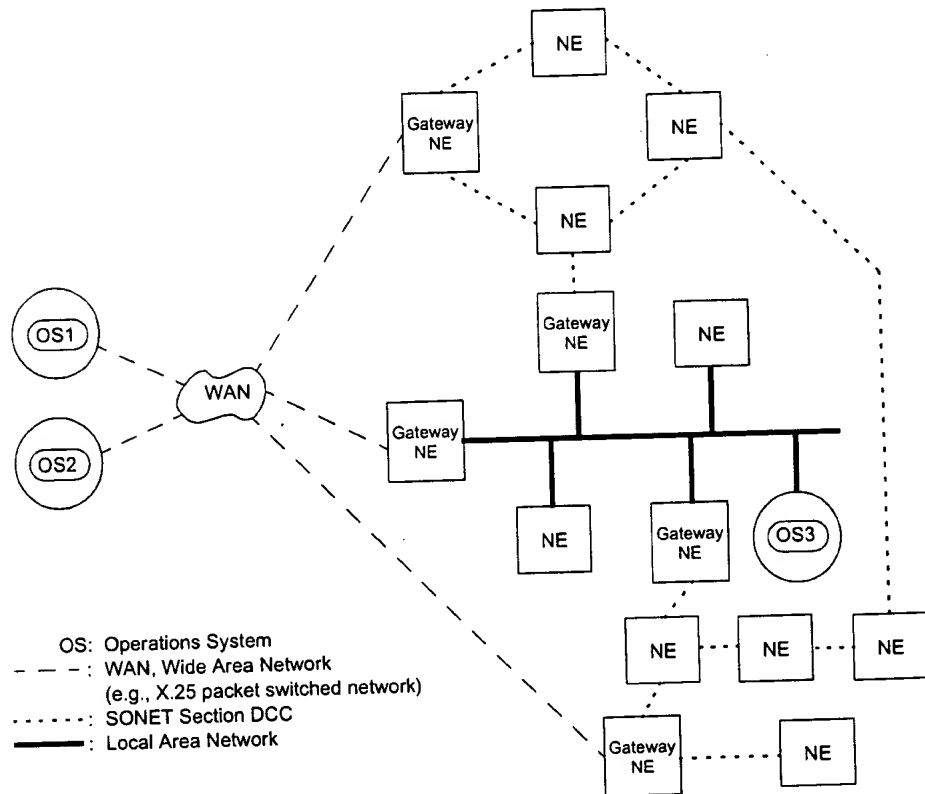
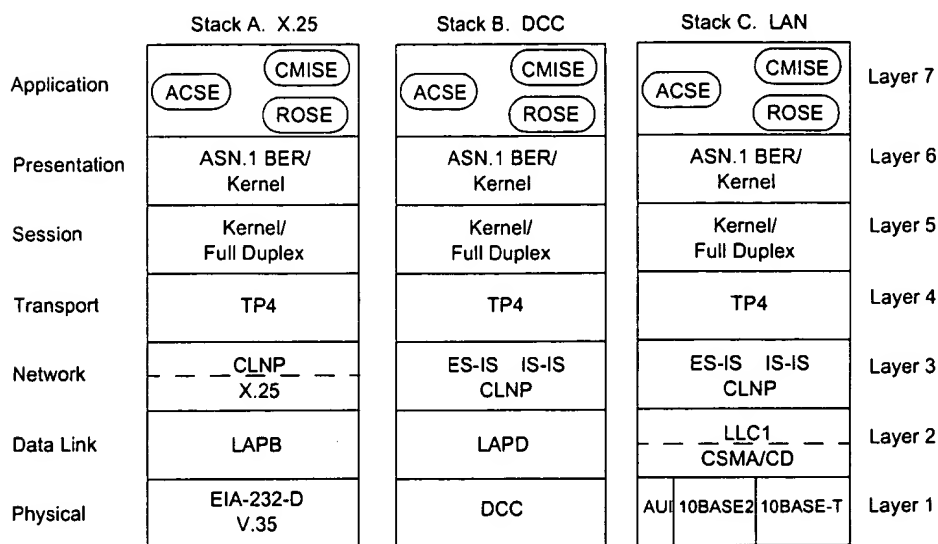


Figure 8-2. Example SONET Operations Communications Architecture

Network providers' initial deployments will likely be much more limited than the example in Figure 8-2, with simpler operations communications networks as shown in Figures 8-4 and 8-5. Figure 8-4 shows two OSs (such as a memory administration OS and a surveillance OS) communicating with a SONET network via a WAN. Most of the SONET NEs are in the Central Office (CO) on a LAN with one far-end SONET NE in a point-to-point DCC configuration. This example also illustrates how mediation devices might be used to provide gateway functions such as interworking the WAN with the LAN. Section 8.1.2 discusses gateway functions, and Section 8.4 discusses interworking.



ACSE: Association Control Service Element
CMISE: Common Management Information Service Element
ROSE: Remote Operations Service Element
ASN.1: Abstract Syntax Notation 1
BER: Basic Encoding Rules
TP4: Transport Class 4
CLNP: Connectionless Network Protocol
LAPD: Link Access Procedure - D Channel
IS-IS: Intermediate System to Intermediate System

DCC: Data Communications Channel
LAPB: Link Access Procedure - B Channel
LLC1: Logical Link Control
CSMA/CD: Carrier Sense Multiple Access with Collision Detection
AUI: Attachment Unit Interface
10BASE2: 10 Mb/s Baseband Coax Cable
10BASE-T: 10 Mb/s Baseband over Twisted Pair
ES-IS: End System to Intermediate System

Figure 8-3. Interactive Protocol Stacks for SONET Operations Communications

Figure 8-5 shows an example of how the SONET DCC may be used in a survivable ring application. In this example, two of the NEs on the ring support gateway functions for added OS-NE operations communications reliability (one gateway is primary and the other provides a backup).

Depending on an NE's placement and application within a network, it may be a Gateway NE (GNE), Intermediate NE (INE), or End NE (ENE). Figure 8-6 shows the same operations communications architecture as in Figure 8-2, but the nodes (NEs) have been labeled by the operations communications role they perform. Note that the OSs are not explicitly labeled with the type of role they may play.

The role that a given SONET NE supports (i.e., GNE, INE, or ENE) is determined by the operations communications network architecture. Thus, network providers should work closely with equipment suppliers to ensure that the operations communications functions

provided by SONET NEs meet the needs of individual architectures. (This may include possible migration strategies to more complex operations communications network architectures.)

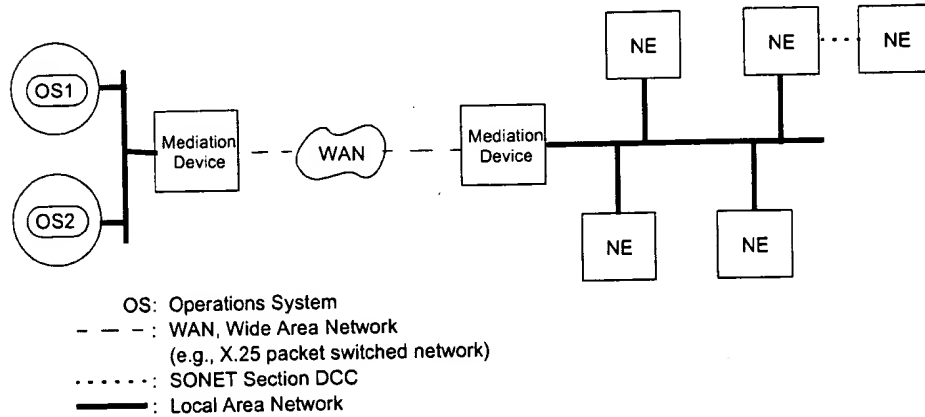


Figure 8-4. Example Intra-site LAN and Point-to-Point DCC

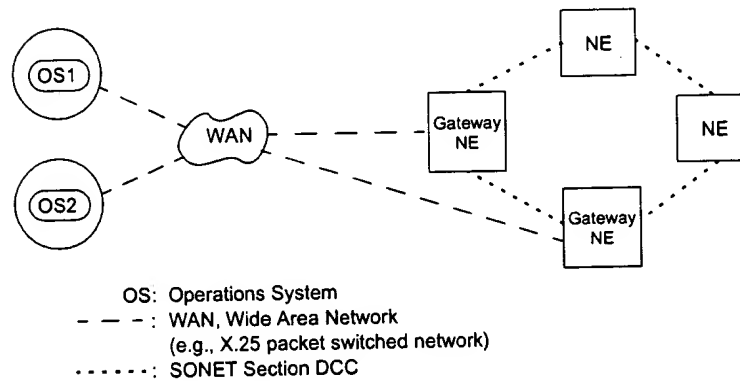


Figure 8-5. Example Operations Communications Network for a Survivable Ring

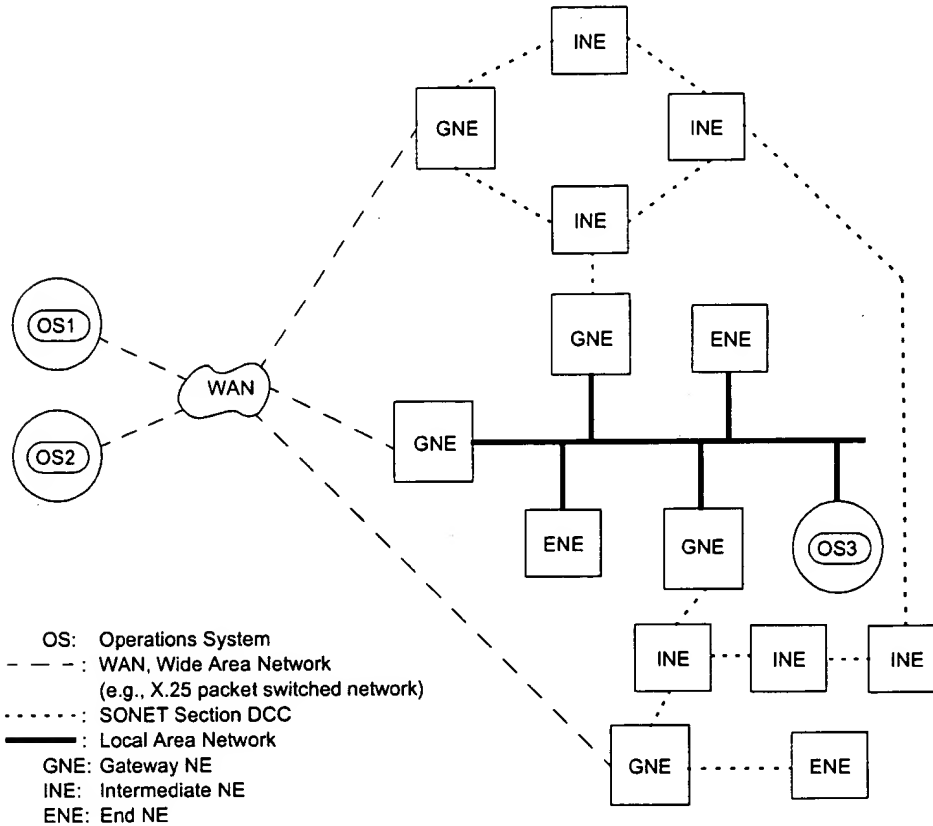


Figure 8-6. Operations Communications Functions

8.1.2 Gateway NE Requirements

A GNE interworks two different kinds of networks. In SONET, there are three possible types of GNEs:

1. A GNE that interworks an X.25 WAN (protocol stack A in Figure 8-3) and the SONET DCC (protocol stack B in Figure 8-3)
2. A GNE that interworks the X.25 WAN and an intra-site LAN (protocol stack C in Figure 8-3)
3. A GNE that interworks an intra-site LAN and the SONET DCC.

This interworking takes place at layer 3, the Network layer, and is transparent to the management operations such as alarm surveillance.

Another GNE function is message concentration for the X.25 WAN. Instead of having one X.25 virtual circuit (VC) to each SONET NE, the gateway can provide an X.25 VC between it and the OS which can be used for messages to and from the OS and other subtending NEs on the network. The DCC or LAN is used to transport the messages within the SONET operations communications network.

In the near-term, before full OSI/CMISE functionality is available in OSs and in SONET NEs, TL1 may be used for operations messages. The OS-NE interface may support TL1 messages over the X.25 protocol, and the NE-NE interface may support TL1 over a seven-layer OSI protocol stack. Section 8.3 contains detailed requirements for NE support of the seven-layer protocol stack on the DCC and the use of CMISE or TL1 messages on this stack. It also contains proposed generic requirements for support of the seven-layer OSI protocol stack on a LAN. With this near-term TL1 configuration, the gateway must perform some application-layer processing of the TL1 messages to interwork the TL1/X.25 protocol stack and the TL1/OSI protocol stack (DCC or LAN). Details of this application-layer processing for TL1/X.25 interworking with the TL1/OSI DCC are provided in Section 8.4. Details of the similar processing for interworking with LANs are also provided in Section 8.4. Section 8.5 describes the ISO routing protocols that can be used to provide selective routing and contains generic criteria for SONET NE support of these protocols.

SONET GNEs must support both level 1 and level 2 Intermediate System-to-Intermediate System (IS-IS) routing protocol (ISO 10589) functions as well as the IS role of the End System-to-Intermediate System (ES-IS) protocol (ISO 9542).

- R8-1** [776] All SONET Gateways shall support level 1 and level 2 IS-IS routing functions as defined in ISO 10589 and in Appendix C.
- R8-2** [777] A SONET Gateway shall support the IS Role of the ES-IS routing protocol over LAN and DCC interfaces as defined in ISO 9542 and Appendix C.
- R8-3** [778] A SONET Gateway shall support the IS-IS Reachable Address Prefix functionality defined in ISO 10589.¹

A detailed description of level 1 and level 2 routing is in Section 8.5. A feature of IS-IS which may be used when both level 1 and level 2 routing is supported is partition repair. If a routing area becomes partitioned because one or more links have failed, level 2 routing (with partition repair) can be used to repair the partition.²

- O8-4** [779] A SONET Gateway should support the level 1 partition repair function as defined in ISO 10589 and in Appendix C.

1. Reachable address prefixes are used to facilitate routing of messages from NEs to OSs over dynamically assigned circuits.

2. Appropriate links must be in place for partition repair to take place.

8.1.3 Intermediate NE Requirements

An Intermediate NE has one or more subtending NEs and performs routing for tandem traffic. An INE must support IS-IS level 1 routing and the IS role of the ES-IS protocol.

- R8-5** [780] A SONET Intermediate NE shall support level 1 IS-IS routing functions as defined in ISO 10589 and in Appendix C.
- R8-6** [781] A SONET Intermediate NE shall support the IS role of the ES-IS protocol over LAN and DCC interfaces as defined in ISO 9542 and Appendix C.

8.1.4 End NE Requirements

End NEs handle only their own local traffic. End NEs must have direct access to either a DCC, a LAN, or an X.25 WAN (to an OS), but they have no subtending network. An example of such a system may be a SONET NE with only one DCC connection (e.g., a terminal multiplex at the end of a feeder route) or a SONET NE on a LAN that is not a GNE. An ENE must support the ES role of the ES-IS protocol.

- R8-7** [782] A SONET End NE shall support the ES role of the ES-IS protocol over LAN and DCC interfaces as defined in ISO 9542 and Appendix C.

8.1.5 Mediation Device (MD)

An MD is a network entity that performs mediation functions. Mediation functions consist of communications gateway functions and possibly additional information processing functions for a particular subnetwork of SONET NEs (see Figure 8-1). Stand-alone MDs might also be used to perform mediation functions in the SONET TMN. Communications gateway functions are essentially the same as those a Gateway NE performs. The decision to use a stand-alone MD rather than a Gateway NE is one that usually depends on the complexity of the communications gateway functions to be performed. For example, if extensive message translation functions that would overburden an NE are required, a stand-alone MD may be warranted. The decision to use an MD (either stand-alone or as an independent module of an NE) to provide gateway communications functions can also be influenced by a supplier's plans to include information processing functions within the MD. Information processing functions are sometimes included in supplier MDs to provide common NE processing functions or subnetwork management functions on a centralized basis for a given subnetwork. Although generic requirements for such capabilities are not discussed in this GR, the possible use of MDs (primarily to support complex communications gateway functions) is included in the discussions below. (See Glossary for further definition of mediation functions.) Additional information regarding Element

Management Layer (EML) applications that may be provided by suppliers in an MD can be found in the following documents:

- TA-TSV-001294, *Generic Requirements for Element Management Layer (EML) Functionality and Architecture*
- FA-NWT-001345, *Framework Generic Requirements for Element Manager (EM) Applications for SONET Subnetworks*
- SR-TSV-002671, *EML Applications for Fault Management: Subnetwork Root Cause Alarm Analysis*
- SR-TSV-002672, *EML Applications for Fault Management: Intelligent Alarm Filtering for SONET*
- SR-TSV-002675, *EML Applications for Configuration Management: Resource Provisioning Selection and Assignment - Functional Description*
- SR-TSV-002678, *EML Applications for Configuration Management: Inventory Notification and Query - Functional Description.*

8.2 Communication Types

The primary types of operations communications that a SONET NE must be able to support are OS/NE, MD/NE, NE/NE and Craftsperson/NE communications, which are each described in the following sections.

8.2.1 OS/NE Communications

OS/NE communications are required at all SONET NEs to perform network operations and management functions. An OS/NE communications path can be direct or indirect. A direct OS/NE communications path is one with no gateway or intermediate systems between the OS and the NE. The direct path may be a dedicated physical connection, an IEEE 802.3 LAN, or through an X.25 WAN. An indirect OS/NE path may consist of at least one NE/NE or MD/NE interface and access to an OS/NE interface via gateway functions. SONET OS/NE communications require the use of a message-oriented channel. In initial deployments and in certain network configurations, network providers may want SONET NEs to be directly connected to an OS (i.e., a direct communications path). For example, a direct connection may be made to each NE for operations communications with an OS through dedicated physical connections. Such NEs would not necessarily be supporting Gateway NE functions even though they support direct OS/NE communications.

R8-8 [783] All SONET NEs shall provide an OS/NE communications path.

R8-9 [784] A SONET NE shall be capable of being equipped with a direct OS/NE interface.

O8-10 [1035] A SONET NE should support the capability to communicate with an OS via a LAN interface.

8.2.2 MD/NE Communications

MDs may aid in various operations communications functions as well as perform operations (information processing) functions for a subnetwork of SONET NEs. When an MD is used in a subnetwork, its primary purpose is to facilitate OS/NE communications and, perhaps, to support remote craftsperson access. To perform these functions, an MD must be able to communicate with the NEs within the subnetwork, OSs, and a local craftsperson terminal. MD/NE communications would include a message-oriented MD/NE interface (see Figure 8-1) and, possibly, one or more message-oriented NE/NE interfaces. The MD would also need to support an OS/MD interface and a WS/MD interface. The language and protocol specifications for these interfaces are identical to the OS/NE and WS/NE interfaces, respectively.

If an MD is used in a SONET subnetwork, it must communicate with each NE in the subnetwork as either a part of OS/NE communications and remote craftsperson communications (because of communications gateway functions of the MD) or for MD/NE communications (because of information processing functions of the MD). An NE may interface directly with an MD via an MD/NE interface, or may do so indirectly via an MD/NE interface and one or more NE/NE interfaces (see Figure 8-1).

If mediation functions are provided within an NE, then the NE/NE interface supported by that NE will be used to interface with other NEs in the subnetwork. If mediation functions are provided in the form of a stand-alone MD, then a LAN is used for the MD/NE interface.

R8-11 [785] When a stand-alone MD is used in a SONET subnetwork, the language and protocol stack for the MD/NE interface shall be identical to that for the NE/NE interface via a LAN.

8.2.3 NE/NE Communications

SONET NEs need to communicate with each other to report alarm, failure, status, and error indications (such as AIS), and to perform protection switching. To support NE/NE communications, a SONET NE is required to have NE/NE interfaces. NE/NE communications are provided in two forms: bit-oriented EOCs and message-oriented operations channels (such as the Section DCC or LAN). The use of the bit-oriented EOC vs. a message-oriented channel depends on the type of information to be communicated and the time-critical nature of the communications (e.g., AIS needs to be transmitted and

received within microseconds and therefore utilizes a bit-oriented EOC). If two SONET NEs need to communicate with each other using a message-oriented channel, the same message-oriented channel used in indirect OS/NE communications will be used to transport this message-oriented NE/NE traffic. A separate message-oriented channel is not needed.

NE/NE message-oriented operations communications will initially be involved primarily as part of indirect OS/NE communications. However, it is expected that NE/NE message-oriented communications to support peer-to-peer operations communications among SONET NEs will become more prevalent as more complicated SONET architectures are implemented and as operating the network requires more sharing of information among the SONET NEs in these architectures.

Note that the choice of utilizing the Section DCC, a LAN, or both for intra-site NE/NE communications is up to the network provider based on their operations communications architecture plans.

R8-12 [786] All SONET NEs shall support NE/NE operations communications paths.

R8-13 [787] To support message-oriented NE/NE operations communications, a SONET NE shall be capable of being equipped with LAN and Section DCC interfaces to other NEs.

8.2.4 Craftsperson/NE Communications

The interfaces defined for craftsperson/NE communications are shown on Figure 8-12. Access to the local NE involves both the craftsperson/WS interface and the WS/NE interface. These interfaces are further defined in Section 8.6.

R8-14 [788] Local craftsperson access by means of a WS is required for all SONET NEs.

8.3 SONET Operations Communications Interface

All of the communications types described in Section 8.2 utilize a common set of 7-layer OSI protocols collectively referred to as the SONET Operations Communications Interface. Since the OS/NE WAN subnetwork technology may differ from the NE/NE DCC or LAN subnetwork technology, the Physical, Data Link, and Network Layers are addressed separately for the OS/NE and for the NE/NE interfaces.³ In particular, Layers 1 to 3 of the OS/NE X.25 interface rely on GR-828-CORE, *OTGR Section 11.2: Generic Operations*

3. For the NE/NE interface, the protocols for the Physical and Data Link Layers are also addressed separately for the DCC case and the LAN case.

Interface – OSI Communications Architecture, for detailed requirements. However, the Transport, Session, Presentation, and Application Layers are defined to be identical for both SONET NEs and OSs supporting SONET NEs.

Most of the criteria for the SONET Operations Communication Interface are based on the criteria in GR-828-CORE. Both the Interactive Protocol Stack and the File-oriented Protocol Stack defined by GR-828-CORE are used for the SONET OS/NE X.25 interface. Additional detail has been supplied in both the text of this GR and in Appendices C and D to further specify the interface and to help assure interoperability both between SONET NEs and between SONET NEs and OSs. For example, this GR and its appendices specify as mandatory for the SONET Operations Communications Interface some protocol capabilities that are optional or out of scope in GR-828-CORE. Appendices C and D provide a protocol profile for the SONET Operations Communications Interface. Appendix C addresses the Data Link Layer (LLC for LANs and LAPD for the DCC), the Network Layer (CLNP, ES-IS, and IS-IS routing protocols), and the Transport Layer. Appendix D addresses the Session and Presentation Layers and ACSE in the Application Layer. In the future, these Appendices may be expanded to include other protocols.

The use of X.500-based Directory Services (see **O8-20 [1036]**) is based on ANSI T1.245, *Directory Service for Telecommunications Management Network (TMN) and Synchronous Optical Network (SONET)*.

A supplier may choose to implement these requirements in a phased approach.

- R8-15** [789] SONET NEs shall provide an appropriate SONET Operations Communications Interface that conforms to the protocol profiles specified in Appendix C, SONET Operations Communications Protocol Profile – Lower Layers, and Appendix D, SONET Operations Communications Protocol Profile – Upper Layers.
- R8-16** [790] SONET NEs shall support CMISE as the application layer protocol for Interactive Class communications.
- R8-17** [791] When file-oriented applications are supported, SONET NEs shall support FTAM as the application layer protocol as specified in GR-1250-CORE.
- CR8-18** [792] SONET NEs may, on an interim basis, support TL1 as the application layer protocol for Interactive Class communications.
- CR8-19** [793v2] The SONET OS/NE X.25 interface may, on an interim basis, support the non-OSI communications architecture (TL1 over X.25) as specified in TR-TSY-000827, *OTGR Section 11.1: Generic Operations Interfaces: Non-OSI Communications Architecture*.

- O8-20** [1036] When a SONET NE supports CMISE on the OS-NE or the NE-NE interface, it should also support the X.500-based Directory Services for TMN and SONET, as defined in ANSI T1.245, for the name/address translation service.

8.3.1 Physical Layer

8.3.1.1 OS/NE

- R8-21** [794v2] At OS/NE X.25 interfaces, SONET NEs shall support the Physical Layer requirements of the TP4/CLNS Protocol Case as described in GR-828-CORE.

At OS/NE-LAN interfaces, the applicable Physical layer criteria are those specified in Section 8.3.1.2 (for NE/NE-LAN interfaces).

8.3.1.2 NE/NE – LAN

- R8-22** [795v2] The Physical layer shall support the following 10 Mb/s baseband Media Dependent interface:

- ... • 10BASE-T per ANSI/IEEE Std. 802.3i-1990, (Supplement to ISO/IEC 8802-3-1990/ANSI/IEEE Std. 802.3-1990) System Configurations for Multi-segment 10 Mb/s Baseband networks (Section 13) and Twisted Pair Medium Attachment Unit (MAU) and Baseband Medium, Type 10BASE-T (Section 14).

- ... The electrical interface and connectors shall be as specified in ISO/IEC 8802-3/ANSI/IEEE 802.3.

- CR8-23** [1011] The Physical layer may also be required to support the following 10 Mb/s baseband Media Dependent interfaces:

- ... a. 10BASE2, as specified in ISO/IEC 8802-3/ANSI/IEEE 802.3
- ... b. The media independent Attachment Unit Interface (AUI) as specified in ISO/IEC 8802-3/ANSI/IEEE 802.3.

8.3.1.3 NE/NE – DCC

- R8-24** [796] The Section DCC, a 192-kb/s channel that is carried in 3 Section overhead bytes of the first STS-1 (i.e., the D1, D2, and D3 bytes) in an

STS-N signal, shall be used as the Physical layer of the message-oriented EOC. The order of transmission is bit 1 of D1 (most significant) through bit 8 of D3 (least significant). Data Link protocols shall transmit bits across this channel by placing them into the next most significant bits.

The EOC that uses the Section DCC for its physical layer is referred to as the *Section EOC*. The following EOC protection criteria apply if linear APS or BLSR protection is provided.

R8-25 [797] Section EOCs shall be protected in the same way as working traffic is protected. The protection switch for the EOC shall follow the protection architecture and mode of operation (e.g., if the traffic is protected bidirectionally, the Section EOC is also protected bidirectionally; if the traffic is protected unidirectionally, then the Section EOC is also protected unidirectionally).

R8-26 [798] A SONET RGTR that accesses the Section EOC shall read the K1 and K2 bytes in both directions to determine when an EOC is being carried with working traffic (i.e., to determine when an EOC is usable).

This protection scheme results in the loss of a Section EOC if a loss of working traffic occurs. Also, diversely routed regenerators are not supported by this scheme. For protected lines, the protection scheme can generally protect the Section EOC very quickly. If protection is not available, the network layer routing information distribution protocols (i.e., End System [ES]-Intermediate System [IS] and IS-IS) may still be used to maintain operations communications connectivity (see Section 8.5). The EOC hardware may also have to be protected against failure by an EOC hardware redundancy feature. Individual NE GRs, TRs, and TAs contain the EOC hardware protection requirements.

The Line DCC is located in the D4 through D12 bytes, which are in the line overhead of the first STS-1 of the STS-N signal. Together, these bytes form one 576-kb/s data channel. Use of the Line DCC is described in ANSI T1.105.04-1995, *Synchronous Optical Network (SONET): Data Communication Channel Protocols and Architectures*.

8.3.2 Data Link Layer

8.3.2.1 OS/NE

R8-27 [799v2] At an OS/NE X.25 interface, SONET NEs shall support the Data Link layer requirements of the TP4/CLNS Protocol Case as described in GR-828-CORE.

At OS/NE-LAN interfaces, the applicable Data Link layer criteria are those specified in Section 8.3.2.2 (for NE/NE-LAN interfaces).

8.3.2.2 NE/NE – LAN

- R8-28** [800] Media Access Control functionality for the LAN shall be as specified in ISO/IEC 8802-3 and ANSI/IEEE 802.3 CSMA/CD specifications.
- R8-29** [801] Logical Link Control functionality for the LAN shall be as specified in ISO/IEC 8802-2/ANSI/IEEE 802.2 LLC Class 1 Type 1 service and as described in Appendix C.
- R8-30** [802] The LSAP value 0111 1111, in which the leftmost bit is the least significant bit, shall be used for the LAN. This value would be represented as 'FE' (hex).

8.3.2.3 NE/NE – DCC

- R8-31** [803] The Data Link layer protocol for the SONET Section DCC shall be based on Link Access Protocol on the D-channel (LAPD) as specified in ITU-T Recommendation Q.921, *ISDN user-network interface - Data link layer specification*, and as described in Appendix C.

Note that the Data Link channel can be in one of two states:

1. Active channel state, where LAPD is sending a frame, an abort sequence, or interframe time fill (contiguous flags between frames), or
2. Idle state, where continuous 1's are sent for at least 15 times.

- R8-32** [804] Both the Unacknowledged Information Transfer Service (UITS) and the Acknowledged Information Transfer Service (AITS) shall be supported. AITS shall be the default mode of operation.

The Globally Unique Network Layer Quality of Service (QoS) parameter is used to select between UITS and AITS (see the Network Layer protocol discussion).

Procedures defined in ITU-T Q.921 for using Service Access Point Identifier (SAPI) and Terminal Endpoint Identifier (TEI) subfields of the LAPD address field for LAPD D-channel applications do not apply to the SONET Section DCC applications. LAPD TEI management procedures for D-channel applications also do not apply to SONET Section DCC applications.

- R8-33** [805] The SAPI value shall be preassigned, and shall be settable locally or remotely by an OS.

- R8-34** [806] The Data Link layer procedures, with the exception of the TEI management procedure, shall follow the rules ITU-T Q.921 specifies.
- R8-35** [807] SAPI value of 62 shall be used for SONET Section DCC applications.⁴
- R8-36** [808] The assignment of user-side/network-side roles (and, hence, the C/R bit value) shall be settable and made before initialization of the data link.
- R8-37** [809] A TEI value of 0 (zero) shall be used for SONET Section DCC applications.

8.3.3 Network Layer

The following requirements apply to both the OS/NE and NE/NE interfaces.

Requirements on the N-SEL portion of the NSAP are provided to allow NEs to differentiate between TP4 PDUs and TARP PDUs (see Section 8.7 for a description of TARP).

- R8-38** [810] When TP4 is being run over CLNP, the N-SEL portion of the NSAP shall be set to an initial value of '1D' (hex).
- R8-39** [811] When TARP is being run over CLNP, the N-SEL portion of the NSAP shall be set to an initial value of 'AF' (hex).
- R8-40** [812] The capability to change the value of the N-SEL when the OSI stack is reinitialized shall be supported.

8.3.3.1 OS/NE

- R8-41** [813v3] At an OS/NE X.25 interface, SONET NEs shall support the Network Layer requirements of the CL-WAN profile (CLNS2) as described in IUT-T Recommendation Q.811 (1997), *Lower Layer protocol profiles for the Q3 and X interfaces*, except that use of the ES-IS protocol shall not be supported.

ES-IS will not be used over the X.25 WAN by either the OSs or the NEs. This differs from the requirements found in ITU-T Q.811.

4. The need to reserve additional SAPI values for specific purposes (e.g., DCC maintenance) is for further study.

- R8-42** [1099] At an OS/NE X.25 interface, SONET NEs shall also support the X.25 Subnetwork Service and Protocol requirements found in Section 5.3.2.4 of GR-828-CORE.

Note: SR-104, *Bellcore Digest of Technical Information*, Vol. 14, Issue 12, December 1997, included a Notice to the Industry which announced a proposed correction to GR-828-CORE which is relevant to the above requirements. The proposed correction is to deprecate **R828-159** in GR-828-CORE and replace it with a new requirement which states: "To operate the Connectionless Network Layer Protocol (CLNP) over X.25 subnetworks, the Subnetwork Dependent Convergence Function defined in ISO 8473-3 ITU-T X.622 shall be used."

At OS/NE-LAN interfaces, the applicable Network layer criteria are those specified in Section 8.3.3.2 (for NE/NE-LAN and DCC interfaces).

8.3.3.2 NE/NE – LAN and DCC

- R8-43** [814] The Network layer protocol for DCCs and LANs shall be CLNP (ISO 8473) as specified in Appendix C.
- R8-44** [815] The Subnetwork Dependent Convergence Function (SND CF), as specified in ISO 8473:1988/Add. 3, and the protocol identification convention described in ISO TR 9577, shall be used for the mapping of the primitives defined for the data link services into the underlying service assumed by the CLNP.
- R8-45** [816] The ISO 8473 Category 3 QoS function shall be supported.

The QoS parameter is used to select either UITS or AITS service in the LAPD Data Link Layer protocol.

- R8-46** [817] The coding of the QoS parameter for the selection of UITS/AITS in the data link shall be as follows:
- ... a. The absence of a QoS parameter shall select AITS.
 - ... b. Bits 7 and 8 set to 1 (Globally Unique QoS) and bit 1 set to 1 shall select AITS.
 - ... c. Bits 7 and 8 set to 1 (Globally Unique QoS) and bit 1 set to 0 shall select UITS.
- CR8-47** [818] Other ISO 8473 Category 3 functions may be required except where prohibited below.

- R8-48** **[819v2]** The following service/protocol parameters shall have the values specified below:
- ... a. Error Reporting (E/R) Flag — As stated in ANSI T1.204, the setting of E/R flag is a local matter. The default value of this flag shall be zero to avoid excessive network traffic that can result during broadcast routing.
 - ... b. Partial Source Routing — Partial source routing shall not be supported because NBSIR 88-3824-1, containing OSI implementation agreements, has identified a defect with the partial source routing option that can cause NPDUs to loop in the network until their lifetime expires.
 - ... c. Inactive Subset — Implementations shall not transmit NPDUs encoded using the ISO 8473:1988 inactive subset. Received NPDUs encoded with the inactive subset shall be discarded.
 - ... d. Segmentation — The non-segmenting subset shall not be used. However, implementations shall be capable of receiving and correctly processing NPDUs that do not contain the segmentation part.
 - ... e. Segmentation Permitted (SP) Flag — Implementations shall not generate data NPDUs without a segmentation part (i.e., the SP flag shall be set to 1 and the segmentation part shall be included).
 - ... f. Lifetime Control — The lifetime parameter shall be used as Paragraph 6.4 of ISO 8473:1988 specifies. This parameter shall have an initial default value of at least three times the network span (number of network entities) or three times the maximum transit delay (divided by 500 ms), whichever is greater, as ISO 8473 specifies. The initial default PDU Lifetime Control shall be 10 seconds.
- O8-49** **[820]** The CLNS Congestion Notification should be used. The default value of '0' should be used when originating NPDUs.
- O8-50** **[821]** The mandatory and optional approaches to congestion avoidance and recovery given in NIST Publication 500-202, Part 4, Section 5.1.2.5 should be used.
- R8-51** **[822]** The destination and source addresses used for SONET applications shall be Network Service Access Point (NSAP), as specified in ISO 8348:1993. The Domain Specific Part (DSP) shall be the ISO DCC format as described in ANSI T1.204-1993 and ANSI X3.216-1992.
-

The NSAP is used in the Network layer to address SONET NEs and their supporting OSs in this OSI environment. The NSAP is divided into two components: the Initial Domain Part (IDP) and the Domain Specific Part (DSP). The IDP is further subdivided into the Authority and Format Identifier (AFI) and the Initial Domain Identifier (IDI). The AFI identifies the IDI format and the DSP syntax. For SONET, the value of the AFI is 39 (decimal), which identifies the ISO DCC (Data Country Code) as the address format and preferred binary encoding for the DSP. The ISO DCC is a three-digit numeric code allocated according to ISO 3166. The IDI portion of the IDP has the value of 840 (decimal) for the United States. ISO/IEC 8348:1993 defines a fixed total length for the IDP of 5 digits, or 2 1/2 octets. Under the preferred binary encoding rules, each digit of the IDP is encoded in Binary Coded Decimal (BCD), where each decimal digit is encoded and transmitted using one semi-octet. Due to the odd number of digits allocated to the IDP, it is necessary to pad the IDP with a semi-octet to ensure an integral number of octets as defined in Section A.5.3 of ISO/IEC 8348:1993. The AFI and ISO DCC IDI define the DSP to be composed of 17 binary octets. The breakdown of the entire NSAP, and the 17 octets of the DSP, is shown in Figure 8-7.

FIELD NAME	IDP (Including Pad)			DSP						
	AFI	IDI	IDI PAD	DFI	ORG	RES	RD	AREA	SYSTEM	SEL
NUMBER of OCTETS	1	1 1/2	1/2	1	3	2	2	2	6	1

IDP:	Initial Domain Part	RES:	Reserved
DSP:	Domain Specific Part	RD:	Routing Domain
AFI:	Authority and Format Identifier	AREA:	Identifier for a Routing Area within a Routing Domain
IDI:	Initial Domain Identifier	SYSTEM:	Routing Entity Identifier for Routing Entity within an NE or OS
DFI:	DSP Format Identifier	SEL:	NSAP Selector
ORG:	Organization Identifier		

Figure 8-7. SONET NSAP Format

The encoding of the IDP, including the IDI Pad semi-octet, is shown in Figure 8-8.

IDP (Including Pad)			
Field Name	AFI	IDI	IDI PAD
Decimal Value	39	840	none
BCD encoding in NSAP	0011 1001	1000 0100 0000	1111

Figure 8-8. IDP Encoding

The DSP values are allocated by the ISO member body or sponsored organization to which the ISO DCC value has been assigned. For the United States, ANSI has been chosen as the organization responsible for the DSP format. ANSI assigns values to the network providers for the ORG fields. The DFI portion of the DSP is 128 (decimal) to identify the SONET DSP format. This is encoded in binary as 1000 0000. The remaining fields of the DSP (also encoded in binary) are assigned by the network provider and the equipment supplier. The DSP fields are used by the IS-IS routing protocol, to provide information about the hierarchical structure of routing areas and domains. The NSAP Selector serves to differentiate multiple entities (for example, TP4 or TARP entities) operating over the same Network entity. Section 8.5 has more information about the routing protocols.

R8-52 [823] The System ID field shall be assigned a 6 octet IEEE address by the equipment supplier.

The other assignable fields are provisioned by the network provider.

8.3.4 Transport Layer – OS/NE and NE/NE

R8-53 [824] Class 4 of ISO 8073 (8073 Add. 2, TP4) shall be supported as specified in Appendix C.

TP Class 4 over CONS and TP Classes 0, 1, 2, and 3 are not used for SONET applications.

R8-54 [825] TP4 implementations shall be capable of receiving and processing all possible parameters for all possible TPDUs, dependent upon the class and optional functions implemented.

R8-55 [826] If the ISO Session Layer is being run over TP4, then the TSAP-ID shall be set to an initial value of "TT" (ASCII), i.e., '5454' (hex).

R8-56 [827] The capability to change the value of the TSAP-ID (within a range of 0 to 4 octets) when the OSI stack is reinitialized shall be supported.

The following Transport layer requirements are included in ISP 10608-1.

R8-57 [828] An unknown parameter in any received CR TPDU shall be ignored.

R8-58 [829] Known parameters with invalid lengths in a CR or CC TPDU shall be handled as a protocol error.

R8-59 [830] Known parameters with valid lengths but invalid values in a CR TPDU shall be handled as follows:

- | | | |
|-----|--------------------------------|-------------------------------------|
| ... | a. TSAP-ID: | Send a Disconnect Request (DR) TPDU |
| ... | b. TPDU size: | Ignore parameter, use default |
| ... | c. Version: | Ignore parameter, use default |
| ... | d. Checksum: | Discard CR TPDU |
| ... | e. Alternate protocol classes: | Protocol error |

R8-60 [831] Unrecognized or inapplicable bits of the additional options parameter shall be ignored.

8.3.5 Session Layer – OS/NE and NE/NE

R8-61 [832] The ISO Session layer shall be supported as specified in Appendix D.

R8-62 [833] If the ISO Presentation Layer is being run over the ISO Session layer, then the Session Selector parameter shall be set to an initial value of "SS" (ASCII), i.e., '5353' (hex).

R8-63 [834] The capability to change the value of the Session Selector parameter when the OSI stack is reinitialized shall be supported.

O8-64 [835] An SS-user-data size of up to 65,535 octets should be supported.

The above objective may become a requirement in the future.

8.3.6 Presentation Layer – OS/NE and NE/NE

R8-65 [836] The ISO Presentation layer shall be supported as specified in Appendix D.

R8-66 [837] The P-SEL shall be no greater than 4 octets in length.

The following Presentation Selector values are used to differentiate between various SONET Application Service Elements (ASEs). These values apply to the called presentation selector that must be included in connect presentation PPDUs (see Appendix D).

R8-67 [838] When CMISE PDUs are sent, the P-SEL shall be set to an initial value of '01' (hex).

R8-68 [1037] When X.500 Directory Access Protocol (DAP) PDUs are sent from the Directory User Agent (DUA) to the Directory System Agent (DSA), the P-SEL shall be set to an initial value of '04' (hex).

R8-69 [839] When FTAM PDUs are sent, the P-SEL shall be set to an initial value of '02' (hex).

R8-70 [841] When TL1 PDUs are sent, the P-SEL shall be set to an initial value of 'AF' (hex).

R8-71 [842] The capability to change the value of the P-SEL when the OSI stack is reinitialized shall be supported.

8.3.7 Application Layer – OS/NE and NE/NE

8.3.7.1 ACSE

R8-72 [843] The Application Control Service Element (ACSE) shall be supported as specified in Appendix D.

The use of the ACSE Authentication Functional Unit is under study. In the future, its use may be required for SONET applications.

8.3.7.2 ROSE/CMISE

SONET NEs shall support CMISE as the application layer protocol for Interactive Class communications (see **R8-16** [790], page 8-12).

R8-73 [844] SONET NEs shall support ROSE/CMISE as specified by GR-828-CORE.

R8-74 [845] The TMN Application Context defined in CCITT M.3100, Section 10, shall be used.

This application context has the same capabilities as the systems management application context defined in ISO/IEC 10040, but also supports the integer values for ProbableCause. The integer value assignments are specified in the CCITT M.3100 ASN.1 Module.

R8-75 [846] The CMISE objects and service mappings for SONET that are contained in GR-1042-CORE and GR-1042-IMD, and the objects and service mappings supporting surveillance and memory administration that are contained in GR-836-CORE and GR-836-IMD shall be supported as per the requirements in those documents.

8.3.7.3 FTAM

When file-oriented applications are supported, SONET NEs must support FTAM as the application layer protocol (see **R8-17 [791]**, page 8-12).

8.3.7.4 Name/Address Translation Services

In order to establish an association between communicating systems over an OSI network, an address, or NSAP, is required. Typically, an OS would know the name of the system it wishes to establish an association with; however, the address of that system is needed. There are a number of ways a translation between the name of a system and the corresponding address can be achieved. A local static mapping table can be used at the OS; in the case of systems named by TIDs (i.e., when TL1 is used) TARP can be used; or an X.500-based Directory Service as defined in ANSI T1.245 can be used. A local static mapping table is a local matter at the OS or NE. TARP is defined in Section 8.7. An objective for the use of X.500 Directory Services is contained in Section 8.3.

8.3.7.5 TL1

SONET NEs may, on an interim basis, support TL1 as the application layer protocol for Interactive Class communications (see **CR8-18 [792]**, page 8-12). The criteria found in this section apply when TL1 is used. Additional requirements for SONET NEs and GNEs to interwork between a TL1-based NE/NE interface and a TL1/X.25 OS/NE interface are provided in Section 8.4.1.

- R8-76** [847] The Presentation context shall contain the TL1 abstract syntax and TL1 transfer syntax that these identifiers specify:

... t11AbstractSyntax OBJECT IDENTIFIER ::= { 1 3 17 104 11 2
bellcoreSONETSyntax(1) }

... t11TransferSyntax OBJECT IDENTIFIER ::= { 1 3 17 104 12 2
bellcoreSONETSyntax(1) }

- R8-77** [848] The transfer syntax for TL1 messages shall be the ASCII encoding of the character string constituting each TL1 message.

- R8-78** [849] The defined context set shall contain the following:

... ABSTRACT SYNTAX { 1 3 17 104 11 2 bellcoreSONETSyntax(1) }
{ joint-iso-ccitt association-control (2), abstract-syntax (1), apdus (0), version
(1) }

... TRANSFER SYNTAX { 1 3 17 104 12 2 bellcoreSONETSyntax(1) }
{ joint-iso-ccitt asn1 (1), basic-encoding (1) }

- R8-79** [850v2] Presentation layer PDUs containing TL1 messages exchanged between communicating SONET NEs shall use the choice of "fully encoded data" for user data defined in CCITT X.226. Also, Presentation data values shall use the "octet-aligned" choice as shown below.

User-data ::= CHOICE {
[APPLICATION 0] IMPLICITSimply-encoded-data,
[APPLICATION 1] IMPLICITFully-encoded-data }

Fully-encoded-data
:= SEQUENCE OF PDV-list

PDV-list ::= SEQUENCE {
Transfer-syntax-name OPTIONAL,
Presentation-context-identifier,
presentation-data-values
CHOICE {
single-ASN1-type [0] ANY,
octet-aligned [1] IMPLICIT OCTET
STRING,
arbitrary [2] IMPLICIT BIT STRING }

Note: When ACSE is used to establish an OSI association for TL1, the "single-ASN1-type" choice (above) is used for ACSE PDU parameters that are mapped to user data parameters in Presentation layer PDUs.

- R8-80** [851] The TL1-ASE shall consist of the appropriate subset of the TL1 messages defined in GR-833-CORE, GR-199-CORE, and GR-834-CORE.
- R8-81** [852] All Application context definitions for associations over which TL1 messages will be exchanged shall include ACSE and the TL1-ASE.
- R8-82** [853] TL1 messages shall be exchanged using Data Presentation Protocol Data Units (TD PPDUs).
- R8-83** [1038] The upper limit on TL1 message size at the application layer, specified in TR-TSY-000827 as 4096 bytes for TL1 over X.25, shall be 4096 bytes for TL1 over any protocol stack or transport mechanism.
- R8-84** [854] Peer NE/NE communications shall use an association established with the Application context identifier below:
- ... t11PeerComm OBJECT IDENTIFIER ::= {1 3 17 104 10 3 t11PeerComm(1)}

See Section 8.4.1 for a discussion of Application contexts for indirect OS/NE communications via the TL1-based NE/NE interface and a GNE.

8.4 Interworking between OSs and SONET NEs

This section provides requirements for SONET NEs and GNEs to interwork the X.25-based OS-NE interface with the DCC or LAN based NE-NE interface for interactive class OS-NE communications. It also discusses interworking the DCC NE-NE interface with the LAN NE-NE interface for interactive communications. This later discussion also applies to interworking the DCC-based NE-NE interface with the LAN-based OS-NE interface.

There are several interworking cases (based on protocols and messages supported) that can be individually examined:

- **TL1/X.25 [OS]–TL1/OSI [SONET]**
 - *SONET LAN Interworking:* TL1 messages are sent between an OS and a SONET NE, using an intra-site LAN (see Figure 8-9).
 - *SONET DCC Interworking:* TL1 messages are sent between an OS and a SONET NE, using a DCC network.
 - *SONET LAN and DCC Interworking:* TL1 messages are sent between an OS and a SONET NE, using a DCC network and an intra-site LAN (see Figure 8-10).
- **TL1/X.25 [OS]–CMISE/OSI [SONET]**
 - *SONET LAN and/or DCC Interworking:* TL1 messages are used by OSs, and CMISE messages are used by NEs.

- **CMISE/OSI [OS]–CMISE/OSI [SONET]**
 - *SONET LAN and/or DCC Interworking:* CMISE messages are used by both OSs and SONET NEs.
- **CMISE (or TL1)/OSI [SONET]–CMISE (or TL1)/OSI [SONET]**
 - *SONET DCC and LAN Interworking:* CMISE or TL1 messages are sent between a SONET NE using a DCC network and a SONET NE using an intra-site LAN. This case also applies to messages sent between a SONET NE using a DCC network and an OS using a LAN.

8.4.1 TL1/X.25 [OS]–TL1/OSI [SONET]

In this interworking scenario, the OS-NE interface is TL1 messages carried over X.25, and the NE-NE interface is TL1 messages carried over the seven-layer Section DCC (or LAN) protocol stack. The ASEs used on the NE-NE interface in this case are ACSE and the TL1 ASE (as described in Section 8.3.7). CMISE and ROSE are not used.

Three OS-NE interworking issues need to be addressed:

1. The GNE must determine the NSAP of the destination NE for OS-to-NE messages
2. The GNE must direct autonomous messages from remote NEs to the appropriate OS(s)
3. Responsibility for establishing connections between OSs, GNEs, and target NEs must be determined.

These issues are each addressed below.

8.4.1.1 Determine Destination NSAP

As TL1 messages are sent from the OS to the GNE, destined for a remote NE, the GNE must determine the destination NE's NSAP address to put in the PDU before it routes the message toward the destination. If the GNE supports the multiplexing of TL1 messages for multiple remote NEs (RNEs) on a single X.25 VC between the OS and the GNE (called "multiple RNEs/VC"), then the GNE has to do a TID-to-NSAP mapping for each TL1 message to determine the NSAP of the destination NE. If the GNE supports the use of one VC between the OS and GNE for each destination NE (called "single RNE/VC"), then the GNE has to do an LCN-to-NSAP mapping. With the multiple RNEs/VC method, fewer X.25 VCs are used; with the single RNE/VC method, the need to extract TIDs from each TL1 message is eliminated.

GNEs that support multiple RNEs/VC or a single RNE/VC have different generic criteria that they need to support. The following requirement is common to both types of GNEs.

- R8-85** [855] The GNE shall maintain static information to map subtending NEs' TIDs to NSAP addresses.

This static information (e.g., a table) can be populated and maintained by TARP (see Section 8.7).

- R8-86** [856] If a GNE supports multiple RNEs/VC, then for an OS-NE message, the GNE shall:

- ... a. receive the TL1 message from an X.25 VC
- ... b. extract the TID information from the message and map the TID to the destination NE's NSAP
- ... c. determine the appropriate association to be used or established
- ... d. establish the association (if necessary)
- ... e. forward the TL1 message over the appropriate association to the destination NE.

- R8-87** [857] If a GNE supports single RNE/VC, then the following requirements apply:

- ... a. the GNE shall maintain dynamic information to map X.25 LCNs to NSAP addresses of subtending NEs
 - ... b. when an X.25 VC is established between an OS and a GNE for communications with a remote NE, the GNE shall
 - listen on that VC for a TL1 command
 - extract the TID from the command
 - map the TID and the X.25 VC LCN to the NSAP
 - establish the association with the remote NE
 - ... c. for an OS-to-NE message, the GNE shall
 - receive the TL1 message from an X.25 VC
 - map the LCN from the X.25 VC over which the TL1 message is received to the destination NE's NSAP
 - determine the appropriate association to be used
 - forward the TL1 message over the appropriate association to the destination NE.
-

8.4.1.2 Directing Autonomous Messages

There are two kinds of TL1 messages that a remote NE can send to an OS – responses to commands, and autonomous messages. Responses to TL1 commands are sent by the remote NE back over the same association over which the request was received. The GNE forwards the response to the correct OS based on the association over which it was received.

R8-88 [858] An NE shall send responses to TL1 commands using the same Application association to the GNE on which the request from the OS was received.

Autonomous messages present a problem because OSs do not have NSAPs (in this interim environment), and autonomous messages do not have TIDs. Thus there is no way that the PDU itself can contain any information that helps the GNE forward it to the correct OS.

Using a modified⁵ multiple Application Context Identifier (ACI) method, one association is established between the GNE and the remote NE for each OS with which the remote NE communicates. Each association between the remote NE and the GNE is established using an ACI that identifies the OS that uses the association. Three ACIs are defined for this purpose: one for a maintenance OS, one for a memory administration OS, and one for a testing OS. GR-199-CORE contains a TL1 message that can be used to establish and maintain the table that maps X.121 addresses to ACIs. If an X.25 VC is established between an OS and a GNE that uses an unknown X.121 address, then the GNE should establish the association with the remote NE using the tl1PeerComm ACI. The GNE maintains information that maps each association to a particular OS, so that when a message (either request-response or autonomous message) is sent from a remote NE to the GNE for forwarding to an OS, the GNE knows which OS should get the message.

R8-89 [859] When a remote NE needs to send an autonomous TL1 message to a particular OS, the NE shall send the message to the GNE over the association established with the appropriate ACI (tl1Maintenance, tl1MemoryAdministration, or tl1Test).

R8-90 [860] To support remote NEs that need to send autonomous messages to a particular OS, the GNE shall support the following ACIs on the NE/NE interface in order to direct autonomous messages to the correct OS.

... tl1Maintenance OBJECT IDENTIFIER ::= {1 3 17 104 10 3 tl1Maintenance(2)}
... tl1MemoryAdministration OBJECT IDENTIFIER ::= {1 3 17 104 10 3
tl1MemoryAdministration(3)}
... tl1Test OBJECT IDENTIFIER ::= {1 3 17 104 10 3 tl1Test(4)}

5. This method is modified from the original multiple ACI method described in TR-NWT-000253, Issue 2.

- R8-91** [861] The GNE shall be capable of establishing associations with subtending NEs for communication between a subtending NE and an OS using the following ACI:

... t11PeerComm OBJECT IDENTIFIER ::= {1 3 17 104 10 3 t11PeerComm(1)}

Note that this ACI is the same one defined in Section 8.3.7.5 for peer TL1 communication on the NE-NE interface. Use of associations established with the t11PeerComm ACI are a local matter.

- R8-92** [862] The GNE shall maintain static information to map the X.25 address of each OS to its related ACI. This static information shall be provisionable.
- R8-93** [863] When an association is established between a GNE and a subtending NE for OS-NE communications, the GNE shall dynamically relate the association with the appropriate OS's X.25 virtual circuit.
- R8-94** [864] For NE-to-OS messages, the GNE shall route the TL1 message over an X.25 VC to the appropriate OS based on the association on which it was received.
- R8-95** [865] When an OS needs to communicate with an NE through a GNE, the GNE shall use or establish an association with the "target" NE using the appropriate ACI (t11Maintenance, t11MemoryAdministration, or t11Test). If the OS' X.121 address is not present as part of the GNE's mapping information, then the GNE shall use the t11PeerComm ACI.

How an NE determines which association to use is a local matter. One example of how this might be done is by having it built into the application in the NE. Another example of how it could be done is through the use of a TL1 message telling the NE what autonomous messages should be sent over which associations.

8.4.1.3 Establishing Connections

Establishing end-to-end OS-NE connectivity in this TL1 interworking environment is a two-step process: the X.25 VC between the OS and the GNE must be established, and the OSI association between the GNE and the target NE must be established. Session establishment is one-way, from the OS through the GNE to the target NE. The communication between the OS and the target NE is two-way.

- R8-96** [866] The OS shall initiate the X.25 VC to the GNE over the non-OSI TL1/X.25 interface.

- R8-97** [867] The GNE shall establish OSI Application associations with remote NEs.
- R8-98** [868] If the X.25 VC between the OS and the GNE goes down, either deliberately or through a communications failure, then the GNE shall take down the association(s) with the remote NE(s) using that X.25 VC.
- ... If multiple RNEs/VC are supported, then associations to several remote NEs will be taken down. If a single RNE/VC is supported, then an association to one remote NE is taken down.

8.4.1.4 SONET LAN Interworking

Figure 8-9 shows an example architecture to illustrate this interworking case. In this example, OS1 and the ENE want to exchange messages. OS1 sends a TL1 message containing the TID of the ENE to the GNE. The GNE maps the TID of the ENE to the ENE's NSAP (using TID-NSAP information provided by TARP) and the GNE maps the NSAP of ENE to ENE's LAN MAC address. This NSAP-to-LAN MAC address mapping was learned by the GNE via the ES-IS routing protocol. The GNE then puts the PDU destined for the ENE onto the LAN. This is done using the Subnetwork Dependent Convergence Function (SND CF) for CLNP and ISO 8802 (LAN protocol) as described in ISO 8473. The ENE takes the PDU with its LAN MAC address off of the LAN and processes the TL1 message.

In the reverse direction, the ENE wants to send a TL1 message to OS1. The ENE puts the PDU on the LAN with the LAN MAC address and the layer 3 NSAP address of the GNE (ES-IS is used to learn the LAN MAC address of the GNE). If this message was a response to a TL1 request by an OS, then the ENE would send the request using the Application association on which the message was received, and the GNE would then forward the message to the destination OS. If this was an autonomous message, then the GNE would forward the message to the destination OS.

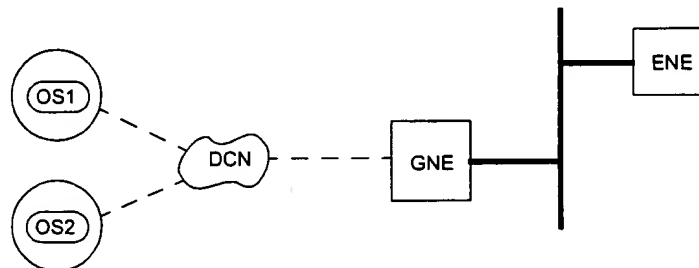


Figure 8-9. TL1/X.25 – LAN Interworking

8.4.1.5 SONET DCC Interworking

In this case, the only interworking that takes place is OS-NE interworking by the GNE. If an OS wants to send a message to a remote NE, it sends a TL1 message containing the TID of the remote NE to the GNE. The GNE maps the TID of the remote NE to the remote NE's NSAP address and then forwards the message along the appropriate DCC toward the destination NE. ES-IS and IS-IS are used to determine on which link to forward the message. Responses or autonomous messages are forwarded from the remote NE to the OS using the modified multiple Application Context Identifier method discussed earlier.

8.4.1.6 SONET LAN and DCC Interworking

Figure 8-10 shows an example architecture to illustrate this interworking case. In this example, OS1 and ENE2 want to exchange TL1 messages. For this to occur, two stages of interworking must take place: X.25-LAN interworking and LAN-DCC interworking. OS-NE interworking using either TID-to-NSAP mapping, or LCN-to-NSAP mapping and multiple Application Context Identifiers, is performed by GNE1.

In addition to this interworking, LAN MAC address mapping must be done. When OS1 sends a TL1 message containing the TL1 TID of ENE2 to GNE1, GNE1 determines ENE2's NSAP, as mentioned. It also maps the NSAP of ENE2 to the LAN MAC address of GNE2; this mapping is learned by the GNE via the IS-IS routing protocol. GNE1 puts the PDU onto the LAN, and GNE2 takes the PDU off of the LAN and routes it on to ENE2.

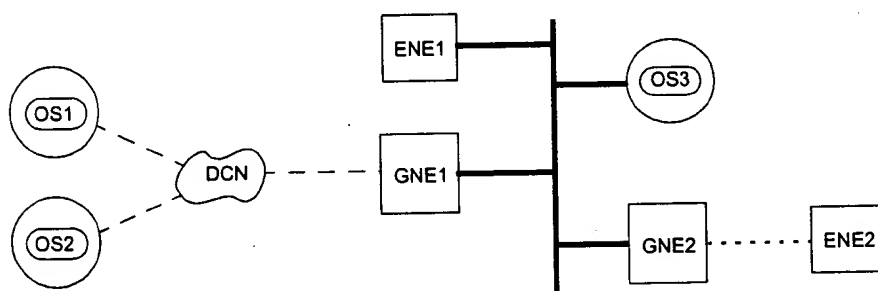


Figure 8-10. TL1/X.25-LAN-DCC Interworking

8.4.2 TL1/X.25 [OS]–CMISE/OSI [SONET]

In the near term environment, the OS-NE interface may be TL1 on a 3-layer X.25 protocol stack. If the NE-NE interface is CMISE on an OSI stack, then the TL1/X.25 protocol stack must be interworked with the CMISE/OSI protocol stack. In addition, the Gateway must perform message translation functions for both OS-to-NE and NE-to-OS messages. These translation functions are assumed to be supplier-specific.

8.4.3 CMISE/OSI [OS]–CMISE/OSI [SONET]

Because SONET NEs are attached to subnetworks that use CLNP (ISO 8473), and OSs are often attached to subnetworks that use the connection-mode network layer protocol (ISO 8208/CCITT X.25), a means to support instances of OS/NE communications must be specified. The most flexible approach to this interworking is included in the methods described in ISO 8648, *Internal organization of the network layer*. This approach corresponds to the Network layer relay operational mode of the interworking functional unit described in ISO TR 10172:1991: *Information technology - Telecommunications and information exchange between systems - Network/transport protocol interworking specification*. The approach is to essentially extend the CLNP out from the SONET DCC, over the ISO 8208/CCITT X.25 subnetwork, into the OS. This requires that the convergence functions as ISO 8473 describes be included in the Gateway NE and in the OS, and that the ISO 8473 protocol be included in the OS. In this method, the ISO 8473 PDUs are embedded in the data field of the ISO 8208/CCITT X.25 packets. Packets arriving at the OS (or Gateway NE) interfaces containing CLNP PDUs are distinguished from others by using the protocol identification conventions described in ISO TR 9577. In the specification of the ISO architecture (ISO/IEC 7498-1:1994, *Information technology - Open System Interconnection Reference Model: The Basic Model*), the view is presented that interworking between subnetworks, in cases such as those described above for OSs and NEs, should take place within the Network layer. The Transport layer, and higher layers, should operate strictly on an end-to-end basis between ESs.

The requirements in Section 8.3 specify the protocol stack required to achieve this interworking.

The above discussion and requirements are valid for both LAN NE-NE interfaces and DCC NE-NE interfaces and combinations of such interfaces (e.g., an OS-NE interface connected to a LAN NE-NE interface, which is then connected to a DCC NE-NE interface) when all of the interfaces support the target 7-layer protocol stacks. The ES-IS (ISO 9542) protocol maps local network-dependent addresses to layer 3 NSAPs and the IS-IS (ISO 10589) protocol determines routes for PDUs as they are transmitted between subnetworks. Specific requirements for support of the routing protocols are provided in Section 8.5.

8.4.4 CMISE (or TL1)/OSI [SONET]–CMISE (or TL1)/OSI [SONET]

Figure 8-10 can be used to illustrate DCC-LAN interworking for CMISE or TL1 NE-NE or OS-NE transaction-oriented messages. In this example, ENE1 and ENE2 want to exchange messages. The following discussion also applies to an exchange of messages between OS3 and ENE2. For TL1 NE-NE messages, Section 8.3.7.5 provides requirements on the use of Application Context Identifiers and encoding of TL1 messages. Since both subnetwork types support CLNP at layer 3, the interworking between these networks is done at layer 3, so it is independent of the application protocol used to exchange messages.

If ENE1 wants to send a message to ENE2, ENE1 puts the PDU on the LAN with the NSAP of ENE2. The LAN MAC address in the PDU will be either the address of GNE1 or GNE2 (ES-IS is used to get the LAN MAC addresses of the GNEs). If GNE2 gets the PDU, it will route the PDU onto the DCC destined for ENE2. If GNE1 gets the PDU, it will send the PDU to GNE2 (using IS-IS to determine that ENE2 is reachable via GNE2), and GNE2 will route it on to ENE2. GNE1 will then send a Redirect message (an ES-IS function) back to ENE1 telling ENE1 that ENE2 is reachable via GNE2 so that the next PDU sent from ENE1 to ENE2 can go directly to GNE2.

If ENE2 wants to send a message to ENE1, it puts the NSAP of ENE1 in the PDU as the destination address and routes it to GNE2 via ES-IS on the DCC. When GNE2 gets the PDU, it puts the LAN MAC address of ENE1 into the PDU (using ES-IS to map ENE1's NSAP to its LAN MAC address) and puts the PDU onto the LAN. ENE1 then takes the PDU off of the LAN and processes it.

Note that the same routing learning process takes place by NEs on a SONET DCC network as on the LAN. ES-IS and IS-IS routing protocols are used by NEs to determine the existence, reachability, and "best routes" to other systems on a DCC subnetwork. IS-IS is also used on a DCC subnetwork to learn which ISs have paths to other subnetworks.

8.5 SONET Operations Communications Routing

8.5.1 Routing Overview

ISO has developed routing protocols (ISO 9542: ES-IS, and ISO 10589: IS-IS) which are to be used by SONET NEs for selective routing of network layer data PDUs. These routing protocols automatically determine the "best" route to all destinations on the network. If link or node failures occur, the ISO 10589 protocol can automatically reconfigure the routing information to "route around" the failure.⁶

ISO has defined an *administrative domain* as a collection of End Systems (ESs), Intermediate Systems (ISs), and subnetworks, operated by a single organization or

6. Note that alternate links must be available for this "routing around" failures to occur.

administrative authority. Network providers will establish their own administrative domains; an example of an administrative domain may be an entire region or a state. A *routing domain* is a set of ESs and ISs that operate according to the same routing procedures and is entirely contained within a single administrative domain. Administrative domains may be broken down into one or more routing domains. Each routing domain can be hierarchically organized into routing subdomains, called *areas*. This is helpful when trying to maintain and process all of the information necessary to perform the routing function by keeping the size of the routing information base and the resources needed to compute routes reasonable. Each routing area maintains detailed routing information about its own internal composition and also maintains information that allows it to reach other routing areas. Since each routing area needs to maintain detailed routing information only about its own internal composition, the amount of data stored in routing information bases is minimized. This, in turn, reduces the amount of data that needs to be exchanged to update this information and the computational overhead associated with computing routes within a routing domain.

Routing within an area is referred to as level 1 routing. Routing between areas is referred to as level 2 routing. Level 2 ISs keep track of the available routes to destination areas. Level 1 ISs keep track of the routing within their own areas. For a Network layer PDU (NPDU) destined to another area, an ES (the source of the PDU) will send the NPDU to an IS in its area. This IS, a level 1 IS, sends the NPDU to the nearest level 2 IS in its own area.⁷ The NPDU then travels via level 2 routing to the destination area, where it again travels via level 1 routing to the destination ES. Figure 8-11 shows an example of a routing domain organized into three areas. An NPDU traveling from ES1 (in Area 1) to ES8 (in Area 3) will go from ES1 to IS1 using ES-IS routing information. IS1 will route the NPDU to IS4 via level 2 routing. IS4 then routes the NPDU to IS 5 via level 1 routing, and IS5 then routes the NPDU to ES8 – the destination system. It is possible for a system to fill several routing roles. In the example, IS1, IS2, IS3, and IS4 are all both level 1 and level 2 routers; IS5 is a level 1 router only. This example shows one way of partitioning a network into areas, and illustrates how level 1 and level 2 routing are used. If this network example was all a single area, then only level 1 routing would be used.

Intradomain routing is facilitated by two different ISO routing information exchange protocols, ES-to-IS (ES-IS) and IS-to-IS (IS-IS). ES-IS routing permits ESs and ISs to exchange configuration and routing information to facilitate routing and relaying. IS-IS routing permits ISs to exchange configuration and routing information to facilitate routing between ISs. In the previous example, illustrated by Figure 8-11, ES1, ES8, IS1, and IS5 support the ES-IS routing protocol. ES1 used it to find out about IS1, and IS5 used it to find out about ES8. All the ISs in that example support the IS-IS routing protocol, which provides both level 1 and level 2 routing.

The NSAP address plays a key role in ISO routing. By examining the NSAP destination address for an NPDU, an IS can determine how to route it. Figure 8-7 shows the structure

7. Since an IS may support both level 1 and level 2 routing, the nearest level 2 IS may be itself.

that is used to specify SONET NE NSAP addresses. A unique area address consists of the area field and all fields to the left of the area field in Figure 8-7 (i.e., the area address includes the AFI, IDI, DFI, ORG, RES, RD, and AREA). The area address is used to determine the next IS to which a PDU should be routed. To determine the next hop for an NPDU, an IS examines the destination NSAP field, which is part of the NPDU. If the area address of the destination NSAP matches its own area address, the IS will use level 1 routing to the next hop. If the area address of the destination NSAP differs from its own area address, then the IS will use level 2 routing to the next hop.

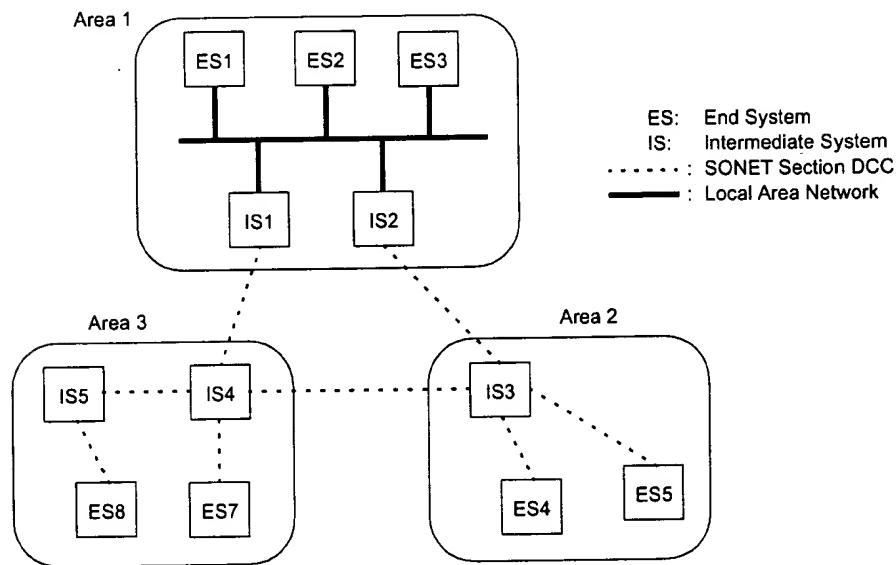


Figure 8-11. Example Routing Domain

8.5.2 ES-IS Requirements

The ES-IS protocol capabilities are organized into two groups: Configuration Information and Redirection Information. Configuration Information is used by ESs to discover the existence and reachability of ISs, and it is used by ISs to discover the existence and reachability of ESs within the same subnetwork. Redirection Information is used by ISs to inform ESs of potentially better routes to use when forwarding NPDUs to a particular destination. Note that since the DCC is point-to-point, Redirection Information is not used, since an ES has only one route to an IS. Redirection Information is used on a LAN. In Figure 8-11, for example, ESs in Area 1 have two ways to send PDUs out of the subnetwork. If ES1 were to send a PDU destined for ES7 (in Area 3) to IS2 for routing on toward its destination, IS2 would route the PDU on toward its destination and then send a

Redirect PDU back to ES1 telling ES1 that IS1 is a better IS to use for sending PDUs to ES7.

- R8-99** [869] All SONET NEs shall support the ES-IS protocol over the NE/NE operations communications interface as specified in Section C.5 of Appendix C.
- O8-100** [870] All SONET NEs supporting the Redirect capability in the ES role of the ES-IS protocol should support the ISO 9542 Address Mask Generation function.
- O8-101** [871] All SONET NEs supporting the Redirect capability in the IS role of the ES-IS protocol should be able to send the ISO 9542 PDU Address Mask field.

The Address Mask appears only in ES-IS Redirect PDUs. The Address Mask parameter indicates that the redirection information applies to a larger population of NSAP addresses than the Destination Address of the Redirect PDU indicates. The Address Mask establishes an equivalence class of NSAP addresses to which the same redirection information applies.

- O8-102** [872] All SONET NEs supporting the ES or the IS role of the ES-IS protocol should be able to send and receive the ISO 9542 PDU Security field.

Note: the manner in which the security field may be used to augment routing security is an area for further study. In the future, the use of the security field may progress from an objective to a requirement.

8.5.3 IS-IS Requirements

- R8-103** [873] All SONET NEs supporting the IS-IS protocol shall do so in accordance with the protocol specifications in ISO 10589 and Appendix C.
- O8-104** [874] SONET NEs supporting the IS-IS protocol should authenticate IS-IS PDUs based on passwords as specified in ISO 10589.

For X.25-based OS-NE communications, the reachable address prefix capability of IS-IS can be used to facilitate routing of messages from an NE to an OS. OSs are assigned NSAP addresses in a different routing area as indicated by the NSAP address prefix. When an NE puts a PDU out on the network destined for an OS, the level 1 routers in the SONET network recognize the address as belonging outside of their routing area and forward the PDU to a level 2 router which forwards it to an OS-NE GNE. The OS-NE GNE maintains a table of NSAP reachable addresses for each OS and the OS's corresponding X.121 address. Using this table, the GNE routes the message on to the destination OS.

- R8-105** [875] A SONET GNE shall maintain static information to map NSAP addresses of OSs to their corresponding X.121 addresses. This static information shall be provisionable locally or remotely.

There is another use of the reachable address prefix capability. An organization may wish to divide its Administrative Domain into a number of separate Routing Domains. An Inter-Domain Routing Protocol (IDRP) would be used to route between such routing domains. ISO is developing a standard IDRP (ISO DIS 10747). The suitability of that protocol for SONET operations communications networks is an area for further study. To facilitate the construction of multi-domain topologies, the provision has been made in IS-IS to enter inter-domain routing information. This information is in the form of a set of Reachable Address Prefixes, which may be manually provisioned, or, in the future, automatically entered by an IDRP.

8.6 Craftsperson/NE Interfaces

The interfaces defined for craftsperson/NE communications are shown in Figure 8-12. Access to the local NE involves both the craftsperson/WS interface and the WS/NE interface (which are collectively referred to as the local craftsperson interface).

Remote Login is defined as the function or capability that enables a craftsperson to log into a remote NE that has one or more intermediate NEs between the remote NE and the local WS (where the craftsperson is physically located). Remote Login would enable the craftsperson to connect to any remote SONET NE interconnected to the local SONET NE via the DCC or a LAN and appear as if the craftsperson were physically present at the remote NE.

The following two references provide implementation requirements for the remote login function.

SIF-002-1996, *Remote Login Implementation Requirements Specification*, Issue 1, contains requirements for an X-protocol-based remote login function over the SONET Operations Communications Network. The remote login functionality is achieved in two steps. First, message traffic for remote login is sent over the DCN to a Network Management System. Next, the Network Management System interacts with the target NE on behalf of the WS user.

SIF-009-1997, *NE-NE Remote Login Implementation Requirements Specifications*, contains implementation requirements for a remote login function between a WS (Craft Interface Terminal) and a remote SONET NE. The DCC and/or LAN portions of the DCN provides the communications paths.

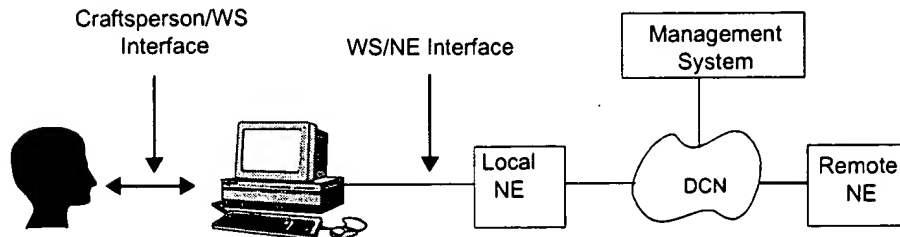


Figure 8-12. Craftsperson/NE Communications Network

8.6.1 Craftsperson/WS Interface

- R8-106** [876] All operations functions supported by a SONET NE via OS/NE communications shall be supported at the local craftsperson interface.
- R8-107** [877] Any features available exclusively at the local craftsperson interface shall be clearly identified in supporting documentation.
- R8-108** [878v2] The workstation shall provide the craftsperson a command-line mode of operation (or interaction) as defined in TR-TSY-000824.
- R8-109** [879] The command-line interface between the craftsperson and the workstation shall conform to TL1 requirements in GR-831-CORE, *OTGR Section 12.1: Operations Application Messages – Language for Operations Application Messages*, which is based on User System Language (USL) of TR-TSY-000825, *OTGR Section 10.A: User System Interface – User System Language*.
- R8-110** [880v2] The user interface requirements specified in GR-826-CORE, *OTGR Section 10.2: User Interface Generic Requirements for Supporting Network Element Operations*, shall be supported (e.g., support of a graphical user interface).

Additional craftsperson/WS interface requirements require further study.

8.6.2 WS/NE Interface

- R8-111 [881v2] A WS/NE interface shall be provided in accordance with TR-TSY-000824.
- O8-112 [1039] It is an objective that SONET NEs support the LAN-based interface defined by Section 4 of SIF-009-1997.
- R8-113 [1040v2] SONET NEs shall support the WS/Remote NE interface defined by Section 5 of SIF-009-1997.

8.7 TARP

The TID Address Resolution Protocol (TARP) is used on the NE-NE interface when there is a need to translate the TID of TL1 messages to the CLNP address (NSAP) of an NE.⁸ Such a need may arise in the following scenarios:

- When TL1/X.25 is used on the OS-NE Interface (see Section 8.4), the Gateway NE needs to be able to map TIDs for subtending NEs to CLNP addresses of those subtending NEs.
- If Remote Login (see Section 8.6) is initiated by entering the TID of the remote NE at the local WS, the local NE needs to be able to map that TID to the CLNP address.

This TID-to-NSAP translation occurs by mapping TIDs to Network Entity Titles (NETs) and then deriving NSAPs from NETs by using the Network Selector Values specified in Section 8.3.3.

TARP uses a selective PDU propagation methodology in conjunction with a distributed database (within NEs) of learned TID/NET mappings. TARP will allow NEs to translate between TID and NET by automatically exchanging mapping information with other NEs without the need for craftsperson intervention. No additional address provisioning is needed at the NE to support TARP.

- R8-114 [882] When a SONET NE supports TL1/OSI on the NE-NE interface (DCC or LAN), the NE shall also support TARP on the NE-NE interface according to the requirements of Sections 8.7 and C.8.
- R8-115 [1041] For all TARP-related TID/NET mappings, case shall be ignored for the TIDs.

A SONET NE may have multiple NETs associated with the same TID. For such NEs, the following requirement applies:

8. TARP is not required to operate in an isolated subnetwork with no INEs present.

- R8-116** [883v2] When an NE supports TARP and has multiple NETs associated with the same TID, the NE shall designate one NET as the NET to be used for mapping purposes for all TARP-related TID/NET mappings.

8.7.1 Network Layer Protocol to Support TARP

The TARP PDU is carried by the standard ISO 8473 (CLNP) Data (DT) PDU. When sending TARP PDUs, TARP places some constraints on the values within the CLNP Data (DT) PDU header fields as specified below. When no TARP constraints are given, these fields will be used according to their specification in ISO 8473.

- R8-117** [884] TARP PDUs shall be carried as ISO 8473 (CLNP) Data (DT) PDUs.
- R8-118** [885] When TARP PDUs are sent, the following constraints shall apply to the header of the CLNP DT PDU:
- ... a. the PDU Lifetime field shall be set to a value of 25000 milliseconds
 - ... b. the Segmentation Permitted flag shall be set to a value of one (1) indicating that segmentation is permitted
 - ... c. the Error Report flag shall be set to a value of zero (0) indicating that discard of the PDU will not cause generation of an Error Report PDU.

8.7.2 TARP PDU Specification

This section specifies the TARP PDU fields that are carried in total by the *Data Part* of the CLNP Data (DT) PDU (see Table 8-1). The following subsections describe each of the TARP PDU fields.

- R8-119** [886] The TARP PDU fields shown in Table 8-1 shall be supported and sent in the order shown by Table 8-1 (starting with tar-lif).

If a node receives a CLNP Service Data Unit (SDU) containing a TARP PDU where the length of the CLNP SDU is greater than N octets [where N is the length of a TARP PDU, which equals 9 octets (the size of the fixed TARP PDU header) + the TID Target Length + TID Originator Length], the SONET NE may consider the TARP PDU invalid and discard it.

Table 8-1. TARP PDU Fields

TARP PDU Fields (within CLNP Data Part)	Abbreviation	Field Size (bytes)
TARP Lifetime	tar-lif	2
TARP Sequence Number	tar-seq	2
Protocol Address Type	tar-pro	1
Update Remote Cache (URC) and TARP Type Code	tar-tcd	1
TID Target Length	tar-tln	1
TID Originator Length	tar-oln	1
Protocol Address Length	tar-pln	1
TID of Target	tar-ttg	n = 0,1,2...
TID of Originator	tar-tor	n = 0,1,2...
Protocol Address of Originator	tar-por	n = 0,1,2...

8.7.2.1 TARP Lifetime (tar-lif)

The tar-lif field contains the TARP time-to-live in hops.

8.7.2.2 TARP Sequence Number (tar-seq)

The tar-seq field contains the TARP sequence number used for loop detection (see Section 8.7.5.7).

8.7.2.3 Protocol Address Type (tar-pro)

The tar-pro field is used to identify the type of protocol address that the TID must be mapped to. The value 'FE' (hex) will be used to identify the CLNP type of address (i.e., NET).

8.7.2.4 URC and TARP Type Code (tar-tcd)

The tar-tcd field consists of the Update Remote Cache (URC) bit (first bit, i.e., most significant bit) and the TARP Type Code (next 7 bits).

The value of the URC bit may be set to 0 or 1.⁹

R8-120 [1042] The URC bit shall be ignored upon receipt of TARP PDUs.

The value of the TARP Type Code identifies the TARP Type of the PDU. Five TARP types are presently defined (see Table 8-2).

Table 8-2. TARP Types

TARP Type	Explanation
1	Request Protocol Address that matches tar-ttg; search Level 1 Routing Area
2	Same as Type 1, but also search Level 2 Routing Area
3	Response to a TARP request
4	Notification of TID or Protocol Address change
5	Request TID that matches Protocol Address (e.g., NET)

If a SONET NE receives a TARP PDU with a TARP Type Code value other than 1 through 5 (i.e., other than the standard values shown in Table 8-2), the SONET NE may consider the TARP PDU invalid and discard it.

8.7.2.5 TID Target Length (tar-tln)

The tar-tln field identifies the number of octets that are present in the tar-ttg field (see Section 8.7.2.8).

8.7.2.6 TID Originator Length (tar-oln)

The tar-oln field identifies the number of octets that are present in the tar-tor field (see Section 8.7.2.9).

8.7.2.7 Protocol Address Length (tar-pln)

The tar-pln field identifies the number of octets that are present in the tar-por field (see Section 8.7.2.10).

9. The URC bit was originally intended to allow the PDU sender to signal the PDU receiver as to whether or not the receiver should update its local cache. Due to security concerns (i.e., the potential for fraudulent use of this feature), it is suggested that the URC feature no longer be used.

8.7.2.8 TID of Target (tar-ttg)

The tar-ttg field contains the TID value for the target NE.

8.7.2.9 TID of Originator (tar-tor)

The tar-tor field contains the TID value of the originator of the TARP PDU.

8.7.2.10 Protocol Address of Originator (tar-por)

The tar-por field contains the protocol address (for the protocol type identified in the tar-pro field) of the originator of the TARP PDU. When the tar-pro field is set to 'FE' (hex) (see Section 8.7.2.3), then tar-por will contain a CLNP address (i.e., the NET).

8.7.3 TARP Data Cache (TDC)

A TARP Data Cache (TDC) may be provided in the SONET NE. When provided, the TDC consists of a set of values for the following triplet: (tar-pro, tar-tor, tar-por). A 4.1 kilobyte TDC could be used to hold approximately 100 entries. For the CLNP case, the TDC is essentially a database of TID-NET mappings.

8.7.4 NE Applications That Use the TARP Processor

Address Resolution and Address Change Notification are two NE applications that will use the TARP processor. These NE applications could be invoked from an Operations System (OS) or from a WS.

- R8-121 [887] The NE shall process address resolution requests (from a higher layer application in the NE) to find the NET that matches a given TID, according to the procedure given in Section 8.7.4.1.
- R8-122 [888] The NE shall process address resolution requests (from a higher layer application in the NE) to find the TID that matches a given NET, according to the procedure given in Section 8.7.4.2.
- R8-123 [889] When a TID or Protocol Address change occurs at an NE, the NE shall notify other NEs of this change according to the procedure given in Section 8.7.4.3.

8.7.4.1 Find NET That Matches TID

When the NE has a TID and needs to find the matching NET, the procedure is the following:

The TARP processor first checks its TDC for the match. If a match is found, the TARP processor would return the result to the requesting application. If no match is found, a TARP Type 1 PDU is originated. If Timer T1 (see Table 8-3) expires, a TARP Type 2 PDU is originated and status information is passed back to the requesting application, indicating that the TARP Type 1 request has failed and that a TARP Type 2 request is being initiated. If Timer T2 expires, then Timer T4 is started, an error recovery routine is initiated, and status information is passed back to the requesting application indicating that error recovery is being initiated.

The error recovery routine is as follows. When Timer T4 expires, another TARP Type 2 PDU is originated and Timer T2 is again started. The *tar-seq* field of this PDU is set to zero, however, the sequence number at the NE is not reset. If Timer T2 again expires, error information is passed back to the requesting application, indicating that the TID could not be resolved.

Table 8-3. TARP Timers

Timer	Description	Default (seconds)	Range (seconds)
T1	Waiting for response to TARP Type 1 request PDU	15	0 - 3600
T2	Waiting for response to TARP Type 2 request PDU	25	0 - 3600
T3	Waiting for response to Address Resolution request	40	0 - 3600
T4	Timer starts when T2 expires (used during Error Recovery)	20	0 - 3600

8.7.4.2 Find TID That Matches NET

When the NE has a NET and needs to find the matching TID, the following procedure takes place:

A TARP Type 5 PDU is originated. Timer T3 (see Table 8-3) is used; however, if this timer expires, no error recovery procedure occurs, and a status message is provided to indicate that the TID could not be found.

A scenario in which this may occur is one in which a Directory Server NE (DSNE) may want to populate its database. A DSNE would typically know which NETs it could

communicate with and could then use TARP to learn the TIDs that correspond to those NETs.

8.7.4.3 Send Notification of TID or Protocol Address Change

When the NE needs to notify other NEs of a TID or Protocol Address change, the procedure is the following:

The TARP Processor originates a TARP Type 4 PDU in which the *tar-ttg* field contains the NE's TID value that existed prior to the change of TID or Protocol Address.

Note that there is no confirmation that other NEs have successfully received the address change information sent in the TARP Type 4 PDU.

8.7.5 TARP PDU Processing

NOTE — The term "originate" is used below to refer to the origination of a TARP PDU from an NE in order to respond to a requesting application within that NE. This term is meant to exclude the propagation of TARP PDUs, which is separately described in Section 8.7.5.8.

TARP PDU processing consists of originating and receiving TARP PDUs.

- R8-124** [890] The NE shall provide the function of a TARP processor that is capable of originating and receiving TARP PDUs for all five TARP Types according to the descriptions given throughout Section 8.7.5.
- R8-125** [891v2] Each time an NE originates a TARP PDU, the NE shall increment the *tar-seq* field. The range of the *tar-seq* field shall be 0 to 65,535. If the *tar-seq* field reaches 65,535 (or if the NE is reset), a TARP Type 4 PDU shall be sent with a *tar-seq* field equal to zero, and the next TARP PDU shall be sent with the *tar-seq* field equal to 1. A zero value will notify all other NEs that a reset has occurred.
- CR8-126** [892] Whenever *tar-seq* is reset to zero, a TARP Type 4 PDU may be generated even if the TID and/or network address has not changed.

Various TARP PDUs must be disseminated to TARP on all neighboring systems. This implies that the NE is capable of identifying its neighbors. The list of neighboring network entities for an NE to transmit to is obtained from the Network Layer Routing Information Base (RIB), and can also include entries created by provisioning.

The set of TARP adjacencies should contain an entry corresponding to each neighbor NET or NSAP in the RIB. (For TARP purposes, the distinction between NETs and NSAPs is unimportant.) When transmitting a TARP PDU to a TARP adjacency, the PDU is sent using the N-UNITDATA Request primitive, using an NSAP constructed by replacing the last octet of the neighbor system's NET or NSAP with the TARP NSEL.

TARP adjacencies can also be provisioned. Note that a TARP adjacency is abstract and need not correspond to a specific data structure maintained by the TARP processor.

The term "adjacency" is used below to mean "TARP adjacency".

8.7.5.1 Origination of a TARP Type 1 PDU

When an NE originates a TARP Type 1 PDU, the PDU is sent to all adjacencies within the NE's routing area. Note that this implies that the NE is capable of identifying its adjacencies. Also note that the NE will typically have more adjacencies when a broadcast subnetwork is used (such as a LAN) than when a point-to-point subnetwork (such as the DCC) is used, and thus would need to send a greater number of PDUs.

- R8-127** [1043] Inclusion of the *tar-tor* field in TARP Type 1 or Type 2 PDUs is optional (at the PDU sender's discretion). The receiver of TARP Type 1 or Type 2 PDUs shall ignore the contents of the *tar-tor* field (if present) and shall be capable of correctly processing the PDUs regardless of whether or not the *tar-tor* field is present.

8.7.5.2 Origination of a TARP Type 2 PDU

When an NE originates a TARP Type 2 PDU, the PDU is sent to all adjacencies both within and outside the NE's routing area within the NE's routing domain. Note that typically only NEs that perform a Level 2 IS function will have adjacencies outside of their routing area. Also note that the use of the *tar-tor* field in a TARP Type 2 PDU may be discussed in a future issue of GR-253-ILR, but that currently that field must be populated with the TID of the Originator (see Section 8.7.5.1).

8.7.5.3 Origination of a TARP Type 3 PDU

A TARP Type 3 PDU is a response to a TARP request PDU. The response is sent only to the originator of the request and thus does not use the TARP propagation procedure (e.g., the receiver of a TARP Type 3 PDU could ignore the *tar-lif* field). The *tar-llg* field of the TARP Type 3 PDU is empty (i.e., zero length).

When an NE has multiple NETs (see R8-116 [883v2]) and is responding to a TARP Type 5 PDU for a NET other than the "designated NET", the *tar-tor* field of the TARP Type 3 PDU is also empty (i.e., zero length).

8.7.5.4 Origination of a TARP Type 4 PDU

A TARP Type 4 PDU is a notification of a TID or Protocol Address change made at the NE that originates the notification (see Section 8.7.4.3). The PDU is sent to all adjacencies both within and outside the NE's routing area.

8.7.5.5 Origination of a TARP Type 5 PDU

When a TARP Type 5 PDU is sent, the CLNP destination address is known and thus the PDU is only sent to that address. Thus TARP Type 5 does not utilize the TARP propagation procedure (e.g., the receiver of a TARP Type 5 PDU could ignore the *tar-lif* field). The *tar-ttg* field of the TARP Type 5 PDU is empty (i.e., zero length).

8.7.5.6 Receipt of a TARP PDU

The following steps are taken by the TARP processor upon receipt on an incoming TARP PDU.

1. Check if *tar-lif* = 0; if so, discard TARP PDU.
2. Check *tar-pro* to see if the Protocol Address Type is supported; if not supported, discard the TARP PDU.
3. Check *tar-seq* and perform the Loop Detection Procedure (only when NE is an IS, see Section 8.7.5.7).
4. The next step depends on the TARP Type Code value and whether the NE is an ES or IS.

8.7.5.6.1 End Systems (ESs)

5. If the TARP Type Code is 1 or 2, then check *tar-ttg*. If *tar-ttg* matches the ES's TID, then originate a TARP Type 3 PDU response.
6. If the TARP Type Code is 3, update TDC and pass response to the requesting application, unless the response is unsolicited and/or a duplicate response in which case the response should be discarded.

7. If the TARP Type Code is 4, then check to see if tar-ttg matches with TDC data. If there is a match, update TDC with the new information.
8. If the TARP Type Code is 5, originate a TARP Type 3 PDU response.
9. If the TARP Type Code is a value that is not supported by the NE, discard the PDU.

8.7.5.6.2 Level 1 Intermediate Systems (ISs)

5. If the TARP Type Code is 1 or 2, check tar-ttg. If tar-ttg matches the IS's TID, originate a TARP Type 3 PDU response. If the tar-ttg does not match the IS's TID, perform Level 1 Propagation (see Section 8.7.5.8).
6. If the TARP Type Code is 3, update TDC and pass response to the requesting application, unless the response is unsolicited and/or a duplicate response in which case the response should be discarded.
7. If the TARP Type Code is 4, check to see if tar-ttg matches with TDC data. If there is a match, update TDC with the new information. In either case, perform Level 1 Propagation (see Section 8.7.5.8).
8. If the TARP Type Code is 5, originate a TARP Type 3 PDU response.
9. If the TARP Type Code is a value other than 1 through 5 inclusive, the PDU may be considered invalid and may be discarded.

8.7.5.6.3 Level 2 Intermediate Systems (ISs)

5. If the TARP Type Code is 1, check tar-ttg. If tar-ttg matches the IS's TID, originate a TARP Type 3 PDU response. If tar-ttg does not match the IS's TID, perform Level 1 Propagation (see Section 8.7.5.8).
 6. If the TARP Type Code is 2, check tar-ttg. If tar-ttg matches the IS's TID, originate a TARP Type 3 PDU response. If tar-ttg does not match the IS's TID, perform Level 1 and Level 2 Propagation (see Section 8.7.5.8).
 7. If the TARP Type Code is 3, update TDC and pass response to the requesting application, unless the response is unsolicited and/or a duplicate response in which case the response should be discarded.
 8. If the TARP Type Code is 4, check to see if tar-ttg matches with TDC data. If there is a match, update TDC with the new information. In either case, perform Level 1 and Level 2 Propagation (see Section 8.7.5.8).
 9. If the TARP Type Code is 5, originate a TARP Type 3 PDU response.
-

10. If the TARP Type Code is a value other than 1 through 5 inclusive, the PDU may be considered invalid and may be discarded.

8.7.5.7 Loop Detection Procedure (performed by ISs)

The Loop Detection Procedure (performed by ISs) is as follows:

Upon receipt of a TARP PDU other than Type 3 or Type 5, the NE checks its Loop Detection Buffer (LDB) for a *tar-por* match. If there is no match, the PDU will be processed and a new couplet entry (*tar-por*, *tar-seq*) is added to the LDB, and if *tar-seq* is zero, a timer associated with the LDB entry is started using the provisionable LDB Entry Timer value. If there is a match, then *tar-seq* is compared to the LDB entry.

If *tar-seq* is non-zero and is \leq LDB entry, the PDU is discarded.

Otherwise, if *tar-seq* $>$ the LDB entry, the PDU is processed and the *tar-seq* field in the LDB entry is updated with the new value. The timer is not affected.

Otherwise, *tar-seq* must be zero. If the LDB entry timer is running, the PDU is discarded. If the timer is not running (i.e., expired), *tar-seq* in the LDB entry remains zero and the associated timer is started as described above.

When the LDB is being populated, only the System ID portion of the *tar-por* address needs to be used. A 4 kilobyte LDB could be used to hold approximately 500 entries. The LDB is flushed periodically in accordance with the LDB Flush Timer.

- R8-128** [893] All NEs supporting an IS function shall maintain a circular (first-in first-out) TARP Loop Detection Buffer.
- R8-129** [1012] All NEs supporting an IS function shall maintain a LDB Entry Timer for each LDB entry for which *tar-seq* = zero. The timer shall be settable within a range of 1 to 10 minutes. The default value shall be 5 minutes.
- R8-130** [894] The LDB Flush Timer shall be settable within a range of 0 to 1440 minutes. The default value shall be 5 minutes.

8.7.5.8 Propagation Procedure (performed by ISs)

The Propagation Procedure (performed by ISs) is as follows:

For *Level 1 Propagation*, PDUs are propagated to all adjacencies within the NE's routing area except as noted below.

For *Level 1 and 2 Propagation*, PDUs are propagated to all adjacencies both within and outside the NEs routing area within the NE's routing domain except as noted below.

The *Propagation Exception* is as follows. On either a point-to-point subnetwork or a broadcast subnetwork, PDUs are not propagated back to the NE from which the PDU was received.

For each NE to which a TARP PDU must be propagated, the NE constructs a new outgoing TARP PDU by decrementing the tar-lif field of the received PDU by one hop and providing new source and destination addresses in the appropriate CLNP header fields. If the decremented lifetime is 0, the system may discard the PDU without further processing, since it would be discarded by any receiving system.

Note that the tar-seq field of the received PDU is not changed during propagation.

The conditions under which TARP PDUs are propagated are given in Section 8.7.5.6.2 (for Level 1 ISs) and in Section 8.7.5.6.3 (for Level 2 ISs). According to those conditions, either Level 1 Propagation or Level 1 and 2 Propagation is performed.

8.7.6 Management of the TARP Processor

At a minimum, SONET NEs shall provide the following capabilities to manage the TARP processor function within the NE.

- R8-131** [895] The NE shall allow TARP propagation to be selectively disabled by link/adjacency.
- R8-132** [896v2] The TARP PDU fields listed in Table 8-4 shall be provisionable in accordance with the default values and ranges specified by Table 8-4. The values of other TARP PDU fields shall not be provisionable:

Table 8-4. Provisionable TARP PDU Fields

TARP PDU Field	Default	Range
tar-lif	100 hops	0 - 65,535
tar-pro	'FE' (hex)	1 byte

- R8-133** [897] The NE shall allow the value of all TARP timers (as shown in Table 8-3) to be provisionable.
- R8-134** [898] The NE shall be capable of displaying (via the local WS at a minimum) the TDC, the LDB, and the TARP Sequence Number in use.

- R8-135** [899] The NE shall provide a manual flush capability for the TDC and the LDB.
- R8-136** [900v2] The NE shall allow manual provisioning of entries for the TDC and the LDB.
- R8-137** [901] The NE shall allow the disabling of any of the following: all TARP functions, TARP propagation functions, TARP origination functions, or the TDC.
- R8-138** [902] The NE shall allow TARP requests to be manually generated.

8.7.7 TARP Echo Function (TEF)

This section describes a TARP Echo Function (TEF) that may be used to aid in troubleshooting and can confirm layer-3 reachability of a CLNP address. The TEF may be invoked from either an OS or from the NE's WS. Invocation of the TEF will result in the NE sending a TARP Type 5 PDU.

When the TEF is invoked, the following information is supplied:

- Address (must be supplied as either a TID¹⁰, a System ID, or a NET)
- How many times to run the TEF (for this invocation)
- When to timeout, i.e., give up waiting for a response to the TEF
- Format for returning results (e.g., results may include round-trip time(s), NET, TID, and TARP Lifetime count).

The response to a TEF invocation returns round-trip time(s) in milliseconds, and success statistics in percent successful. When the TEF is run more than once (for a given invocation), then round-trip times are returned as both individual times and the aggregate time.

Note that if multiple TARP PDUs are outstanding, it is possible to mistake the response from a Type 1 or Type 2 TARP PDU as the response from the TEF. This could result in incorrect TEF information. For this reason, it is suggested that the number of outstanding TARP PDUs at any given time be limited to one.

¹⁰ If the address is supplied as a TID, the address resolution process described in Section 8.7.4.1 would be performed before sending the TARP Type 5 PDU for TEF.

8.7.8 Manual TARP Adjacencies

The use of non-SONET NEs without TARP capability, e.g., generic routers, could cause compatibility issues relating to TARP. Such devices may not have TIDs; however, TARP requests might need to cross a generic router. In such cases, the ability to provision a manual TARP adjacency in the SONET NE may be useful. This manual adjacency would in a sense cause a TARP request to hop through a generic router. This is depicted in Figure 8-13.

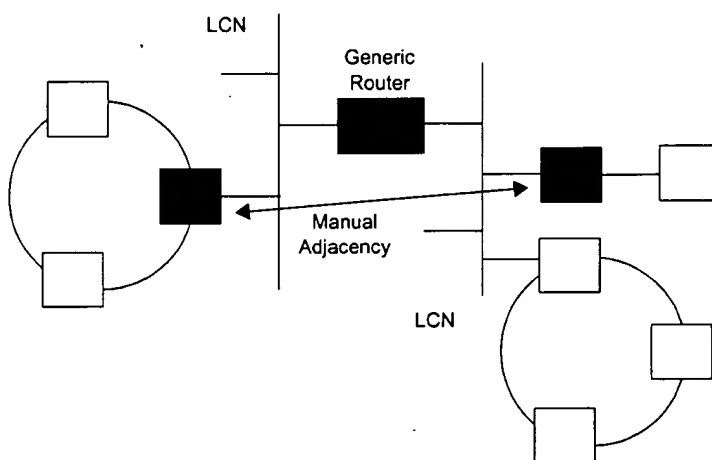


Figure 8-13. Manual TARP Adjacencies

8.7.9 TARP Example

Figure 8-14 illustrates an example of how TARP works. In this example, a Remote Login session is being initiated from the NE with TID AA to the NE with TID HH. The NE with TID AA originates a TARP Type 1 Request PDU. This request is propagated through the network until it reaches the NE with TID HH. At this point, the NE with TID HH originates a TARP Type 3 Response PDU, which is sent back to the NE with TID AA.

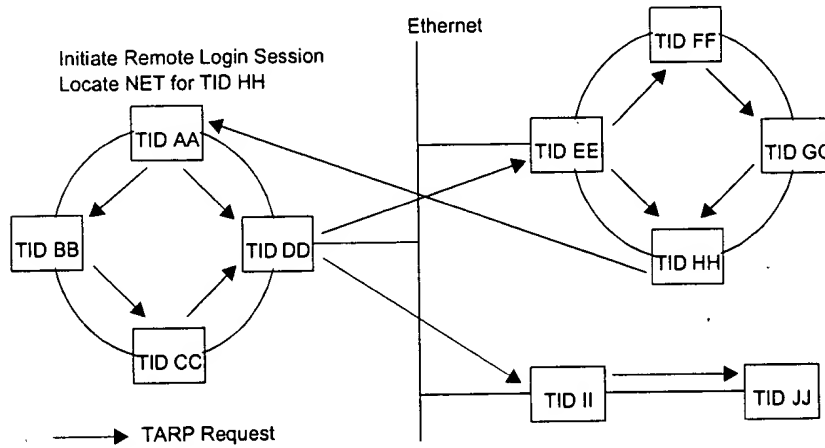


Figure 8-14. TARP Example

8.7.10 TARP Pseudocode

This section provides pseudocode as an aid to the reader in understanding the steps taken by the TARP processor upon receipt of an incoming TARP PDU. The normative description of these steps is provided in Section 8.7.5.6.

BEGIN Pseudocode ALL

End System

Procedure EndSystem()

BEGIN

Check the TARP Lifetime

IF Lifetime has expired THEN Discard packet END PROC

ELSE (*Check to see if the protocol type is supported*)

IF Protocol Type is not supported THEN Discard Packet END PROC

ELSE (*Perform Type Analysis*)

CASE

TARP type is 1 and Target TID is my TID:

Construct Response (Type 3) hand-off to Forward Process END PROC

TARP type is 2 and Target TID is my TID:

Construct Response (Type 3) hand-off to Forward Process END PROC

TARP type is 3:

Add triplet to my cache and pass response to requesting application

Discard duplicate responses END PROC

TARP type is 4:

IF My cache has a TID match THEN

Update my cache with new information

Discard packet END PROC

ELSE My cache does not have a match

Discard packet END PROC

TARP type is 5:

Construct Response (Type 3) hand-off to

Forward Process END PROC

TARP type is not supported or Target TID is not my TID:
Discard packet **END PROC**

END;
END EndSystem;

Level 1 NE

Procedure Propagation(LEVEL)

LEVEL: (Level1, Level2, All)

BEGIN

Consult adjacency and routing database for adjacency
information. Decrement TARP packet lifetime by 1 hop.

IF tar-lif > 0 **THEN**

Construct packet(s) with new destination and source addresses
(PDU Header only, all adjacencies except adjacencies to NE
from which packet was received) and hand-off to Forwarding
process.

END;

END Propagation()

Procedure Loop Detection()

BEGIN (*Loop Check*)

IF Type 3 or 5 **THEN RETURN END PROC**

IF tar-por is a match **THEN** (*Check sequence number*)

BEGIN

IF pdu.tar-seq is non-zero and \leq ldb.tar-seq

OR pdu.tar-seq = 0 and LDB Entry Timer is running **THEN** discard pdu

ELSE IF pdu.tar-seq = 0 and LDB Entry Timer is not running **THEN**

ldb.tar-seq = 0

start LDB Entry Timer

ELSE pdu.tar-seq > ldb.tar-seq

ldb.tar-seq = pdu.tar-seq

END

ELSE tar-por is not a match.

Add couplet (tar-por, tar-seq) to LDB.

IF tar-seq = 0 **THEN** start LDB Entry Timer

END;

END Loop Detection()

Procedure Intermediate System1()

BEGIN

(*Check the TARP Lifetime and Run (Loop Detection Procedure)*)

IF Lifetime has expired and or LDB is a match **THEN** Discard packet

END PROC

ELSE (*Check to see if the protocol type is supported*)

IF Protocol Type is not supported **THEN** Discard Packet

END PROC

ELSE (*Perform Type Analysis*)

CASE

TARP type is 1 and Target TID is my TID:

Construct Response (Type 3) hand-off
to Forward Process **END PROC**

TARP type is 2 and Target TID is my TID:

Construct Response (Type 3) hand-off
to Forward Process **END PROC**

TARP type is 3:

Add triplet to my cache and pass response to requesting
application. Discard duplicate responses **END PROC**

TARP type is 4:

IF My cache has a match **THEN**

Update my cache with new information

BEGIN Propagation(All)

END Propagation **END PROC**

ELSE My cache does not have a match

Do not update cache

BEGIN Propagation(All)

END Propagation **END PROC**

TARP type is 5:

Construct Response (Type 3)
hand-off to Forward Process **END PROC**

TARP type is not supported or Target TID is not my TID:
BEGIN Propagation(Level1)
END Propagation **END PROC**

END;
END IntermediateSystem1;

Level 2 NEs

Procedure IntermediateSystem2()

BEGIN

Check the TARP Lifetime and Run (Loop Detection Procedure)

IF Lifetime has expired and or LDB is a match **THEN** Discard packet
END PROC

ELSE (*Check to see if the protocol type is supported*)

IF Protocol Type is not supported **THEN** Discard Packet

END PROC

ELSE (*Perform Type Analysis*)

CASE

TARP type is 1:

IF TID is my TID **THEN**

Construct Response (Type 3) hand-off
to Forward Process **END PROC**

ELSE Target TID does not match

BEGIN Propagation(Level1)
END Propagation **END PROC**

END;

TARP type is 2:

IF TID is my TID

THEN

Construct Response (Type 3) hand-off
to Forward Process **END PROC**

ELSE TID is not My TID

BEGIN Propagation(All)
END Propagation **END PROC**

END;

TARP type is 3:

Add triplet to my cache and pass
response to requesting application
Discard duplicate responses **END PROC**

TARP type is 4:

IF My cache has a TID match **THEN**
 Update my cache with new information
 BEGIN Propagation(All)
 END Propagation **END PROC**
ELSE My cache does not have a TID match
 Do not update cache
 BEGIN Propagation(All)
 END Propagation **END PROC**
END;

TARP type is 5:

Construct Response (Type 3) hand-off to
Forward Process **END PROC**

TARP type is not supported:

BEGIN Propagation(All)
END Propagation **END PROC**

END;

END;

END IntermediateSystem2;

END Pseudocode ALL

Appendix A: Requirement-Object List

- CR2-1** [1] SONET NEs may be required to provide electrical cross-connect facilities with patching and monitoring jacks for restoration and rearrangements. Page 2-3
- CR2-2** [2] A bridging repeater or equivalent may be required to allow rearrangement from the monitor jack. Page 2-3
- R2-3** [3] Suppliers shall provide a description of their optical fiber distributing frame or optical DSXs. This description shall include the number of terminations, storage provisions for excess fiber, and the type of fiber cable connectors provided (see Section 4). Page 2-4
- R3-1** [4] A SONET NE shall have the capability to ignore the values contained in all undefined and unused bits and bytes [except for Bit Interleaved Parity (BIP)-8 calculations] to prevent misinterpretation of the received patterns. Page 3-2
- O3-2** [5] A SONET NE should send all-zeros patterns (before scrambling) in undefined bits and bytes. All-zeros patterns should also be sent in defined bits and bytes if the NE does not support the defined function or if the function has been disabled by the user. Page 3-2
- R3-3** [6] If a supplier introduces a nonstandard feature employing SONET overhead, the supplier shall disclose such use of overhead, and furnish the network provider with an equipment option to disable the feature (including the transmission of the nonstandard messages). Page 3-3
- R3-4** [7] The structure of an STS-1 shall be as shown in Figure 3-1. Page 3-3
- R3-5** [8] In each byte of the STS-1, the most-significant bit shall be transmitted first, as shown in Figure 3-2. Page 3-4
- R3-6** [9] The structure of an STS-1 SPE shall be as shown in Figure 3-4. Page 3-4

-
- R3-7** [10] The values used to stuff columns 30 and 59 of each STS-1 SPE shall produce even parity in the calculation of the STS-1 Path BIP-8 (see Section 3.3.2.3).
Page 3-5
- R3-8** [11] If an NE supports the multiplexing, switching, or transport of STS-Nc SPEs, then it shall treat each STS-Nc SPE as a single entity.
Page 3-7
- R3-9** [12] The structure of an STS-Nc SPE shall be as shown in Figure 3-7.
Page 3-8
- R3-10** [13] The structure of a VT-structured STS-1 SPE shall be consistent with the structures shown in Figures 3-9 through 3-19.
Page 3-10
- R3-11** [14] Four consecutive 125- μ s frames of the VT-structured STS-1 SPE shall be organized into a 500- μ s superframe, the phase of which is indicated by the H4 (Indicator) byte in the STS POH (see Section 3.4.1).
Page 3-10
- R3-12** [15] The structure of a VT SPE shall be as shown in Figure 3-21.
Page 3-10
- R3-13** [16] The A1 byte shall be set to '11110110' and the A2 byte shall be set to '00101000' in all STS-1s within an STS-N.
Page 3-29
- O3-14** [17] STE that supports line-side signals should have the capability to access the J0 byte, which is located in the first STS-1 of an STS-N.
Page 3-29
- R3-15** [18] Unless it is being used for a defined purpose (e.g., to carry a Section Trace message once the details of that feature are defined) each J0 and Z0 byte shall be set to a binary number corresponding to its order of appearance in the STS-N frame (i.e., the J0 byte shall be set to 00000001, the first Z0 byte shall be set to 00000010, the second Z0 byte to 00000011, etc.).
Page 3-29
- R3-16** [19] The B1 byte in a line-side signal shall carry a BIP-8 code, using even parity. The Section BIP-8 shall be calculated over all bits of the previous
-

STS-N frame after scrambling and placed in the B1 byte of the current STS-N frame before scrambling.

Page 3-30

- R3-17** [20] The B2 byte shall be provided in all STS-1s within an STS-N to carry a Line BIP-8 code, using even parity. The Line BIP-8 shall be calculated over all bits of the Line Overhead and the Envelope Capacity of the previous STS-1 frame before scrambling, and placed in the B2 byte of the current STS-1 frame before scrambling.

Page 3-31

- R3-18** [21] LTE terminating an OC-1 or STS-1 electrical signal shall set bits 5 through 8 of the M0 byte to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected based on the Line BIP-8 (B2) byte. The error count shall be a binary number from zero (i.e., 0000) to 8 (i.e., 1000). The remaining seven values represented by the four REI-L bits (i.e., 1001 through 1111) shall not be transmitted, and shall be interpreted by receiving LTE as zero errors.

Page 3-31

- R3-19** [22] LTE terminating an OC-N or STS-N electrical signal ($N \geq 3$) shall set the M1 byte to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected using the Line BIP-8 (B2) bytes. For values of N below 48, the error count shall be a binary number from zero to 8N. The remaining possible values [i.e., $255 - (8 \times N)$] represented by the eight REI-L bits shall not be transmitted, and shall be interpreted by the receiving LTE as zero errors. For N equal to 48, the count shall be truncated at 255.

Page 3-32

- R3-20** [23] If no message has been loaded by the user for transmission in the J1 byte, then that byte shall be set to all-zeros (i.e., to ASCII NULL characters).

Page 3-33

- R3-21** [24] The B3 byte shall carry a BIP-8 code, using even parity. The STS Path BIP-8 shall be calculated over all bits (783 bytes for an STS-1 SPE or $N \times 783$ bytes for an STS-Nc SPE, regardless of any pointer adjustments) of the previous STS SPE before scrambling, and placed in the B3 byte of the current STS SPE before scrambling.

Page 3-33

- R3-22** [25] For an STS path connection that is equipped and provisioned, a valid non-zero STS Path Signal Label shall be generated by the STS PTE. If the

content of the STS SPE is one of the specific possibilities listed in Table 3-2, then the corresponding code from Table 3-2 (or if PDI-P is supported and one or more Payload Defects are present, Table 3-3) shall be used. If the content is not specifically listed, then the code for "Equipped – Nonspecific Payload" shall be used.

Page 3-34

- R3-23** [26] For STS path connections that are not equipped, or that are equipped but not provisioned, the NE (e.g., the NE's LTE) shall generate all-zeros STS SPEs with "valid" STS Payload Pointers.

Page 3-34

- R3-24** [27] STS PTE shall set bits 1 through 4 of the G1 byte to indicate (to the upstream STS PTE) the count of interleaved-bit block errors that it has detected based on the STS Path BIP-8 byte (B3). The error count shall be a binary number from zero (i.e., 0000) to 8 (i.e., 1000). The remaining seven values represented by the four REI-P bits (i.e., 1001 through 1111) shall not be transmitted, and shall be interpreted by receiving STS PTE as zero errors.

Page 3-36

- R3-25** [28] Bit 1 of the V5 byte shall be set so that the parity of all of the odd-numbered bits (i.e., bits 1, 3, 5, and 7) in all bytes in the previous VT SPE is even. Bit 2 shall be set so that the parity of all of the even-numbered bits (2, 4, 6, and 8) in all bytes in the previous VT SPE is even.

Page 3-39

- R3-26** [29] VT PTE shall set bit 3 of the V5 byte to '1' if one or more errors were detected using the BIP-2. It shall set bit 3 to '0' if zero errors were detected.

Page 3-39

- R3-27** [30v2] For a VT path connection that is equipped and provisioned, a valid, non-zero VT Signal Label shall be generated by the VT PTE. If the mapping contained in the VT SPE is one of the specific possibilities listed in Table 3-4, then the corresponding code from the table shall be used. If the mapping is not specifically listed, then the code for "Equipped – Nonspecific Payload" shall be used.

Page 3-39

- R3-28** [31] For VT path connections that are not equipped, or that are equipped but not provisioned, the NE (e.g., the NE's STS PTE) shall generate all-zeros VT SPEs with "valid" VT Payload Pointers.

Page 3-39

- R3-29** [32] The H4 byte shall be used to indicate phase of the V1 through V4 bytes in the 500- μ s (i.e., 4-frame) VT Superframe. The allocation of the bits in the H4 byte, and the correspondence of the H4 code with the V1 through V4 bytes shall be as shown in Figure 3-26. Page 3-41
- R3-30** [33] If the byte-synchronous DS1 mapping is provided, it shall be as shown in Figure 3-27. Page 3-42
- R3-31** [34] If the P-bits are being used to indicate the phase of the S-bits or the F-bits, then they shall set to the 24-frame sequence shown in Figure 3-28. Page 3-43
- R3-32** [35] VT PTE with byte-synchronous DS1 interfaces shall be capable of accepting DS1 signals using the DS1 Superframe and ESF formats defined in GR-499-CORE. Page 3-45
- R3-33** [36] If the DS1 signal uses the ESF format and the F-bit is not used to transport the framing bits, then the following apply:
- ... • The Cyclic Redundancy Check-6 (CRC-6) code in the received DS1 signals shall be monitored, and detected errors subsequently reported in the Data Link performance report message on the outgoing signal.
 - ... • The correct CRC-6 code shall be calculated and inserted on the outgoing DS1 signals.
 - ... • The NE shall send a performance report message every second to the sink, and receive a performance report message every second from the source.
- Page 3-45
- O3-34** [37] If the NE does not provide DS0 rearrangement capabilities for an incoming DS1, then it should recover clock from the DS1 and use that clock in the creation of the VT SPE. Page 3-45
- CR3-35** [38] The VT PTE may be required to support the signaling transfer mode. Page 3-46
- R3-36** [39] If the signaling transfer mode is being used, then the following apply:

- ...
 - The signaling information carried in the robbed bit positions of the DS1 signal shall be copied to the corresponding S-bit positions within the VT1.5 when the DS1 is mapped. In the VT1.5 to DS1 direction, the signaling information carried by the S-bits within the VT1.5 shall be written over the appropriate robbed bit positions of the outgoing DS1 bit stream.
- ...
 - For DS1s using the Superframe format, the robbed bit positions (from which the signaling information is copied) shall be set to '1' when the DS1 is mapped into the VT1.5.
- ...
 - The phase of the S-bits shall be indicated by the P-bits in the same byte as shown in Figure 3-28 for 2-, 4-, and 16-state signaling schemes.
- ...
 - The VT PTE shall generate a new framing bit pattern for the outgoing DS1 bit stream.

Page 3-46

- CR3-37** [40] The VT PTE may be required to support the clear mode.

Page 3-47

- CR3-38** [41] If the clear mode is being used, then the VT PTE may be required to be user-provisionable (on a per-DS1 basis) to transport the DS1 frame bit in the F-bit position of the VT1.5.

Page 3-47

- R3-39** [42] If the clear mode is being used and the DS1 frame bit is carried in the F-bit, then the framing bits in the incoming DS1 shall be placed in the transmitted F-bits, and the received F-bits shall be placed in the outgoing DS1 signal (with no change in phase relative to the DS0 channels). The phase of the F-bit shall be indicated by the P-bits in the same byte as shown in Figure 3-28 for the Superframe and ESF formats.

Page 3-47

- R3-40** [43] If the clear mode is being used and the DS1 frame bit is not carried in the F-bit (i.e., the F-bit is undefined), then the VT PTE shall generate a new framing bit pattern for the outgoing DS1 bit stream.

Page 3-47

- CR3-41** [44] The VT PTE may be required to be user-provisionable on a per-DS0 basis to either perform or not perform signaling transfer (i.e., to provide signaling transfer/clear DS0 transport).

Page 3-47

-
- CR3-42** [45] The VT PTE may be required to be user-provisionable (on either a per-DS0 or per-DS1 basis) to perform dynamic signaling transfer/clear DS0 transport.
Page 3-48
- R3-43** [46] If dynamic signaling transfer/clear DS0 transport is being used, then the code ABCD=1001 on the S-bits shall be interpreted to mean that signaling transfer is not to be performed on that DS0 channel in either direction of transmission. Signaling transfer shall be performed for DS0 channels whose associated S-bits do not contain the code ABCD=1001.
Page 3-48
- R3-44** [47] The VT PTE shall accommodate both the Alternate Mark Inversion (AMI) line code, and the Bipolar with Eight Zero Substitution (B8ZS) line code.
Page 3-48
- CR3-45** [48] VT PTE that supports DS0 path terminations or DS0 rearrangement capabilities may be required to support ZBTISI for pulse density assurance.
Page 3-49
- CR3-46** [49] VT PTE that supports DS0 path terminations or DS0 rearrangement capabilities may be required to support ZCS for pulse density assurance.
Page 3-49
- CR3-47** [50] VT PTE may be required to support ZBTISI.
Page 3-49
- R3-48** [51] If ZBTISI is supported, then the ZBTISI algorithm and the ESF data link as described in GR-499-CORE shall be used.
Page 3-49
- R3-49** [52] The choice between AMI or B8ZS, and of ZBTISI or ZCS (if provided), shall be provisionable by the user on a per-DS1 interface basis.
Page 3-49
- R3-50** [53] The asynchronous mapping of a DS1 into a VT1.5 SPE shall be as shown in Figure 3-29.
Page 3-49
- R3-51** [54] In each VT1.5 SPE, two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall
-

be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.

Page 3-49

- R3-52** [55] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single bit errors in the C-bits.

Page 3-50

- R3-53** [56] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS1 jitter transfer mask shown in Figure 5-26, the output jitter is less than 0.7 Unit Intervals peak-to-peak (UI_{pp}), assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.

Page 3-50

- R3-54** [903] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS1 jitter transfer mask shown in Figure 5-26, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than that same DS1 jitter transfer mask.

Page 3-50

- R3-55** [57] The DS1 interface shall accommodate both the AMI line code (assuming the DS1 source meets the zeros constraints in GR-499-CORE, see Section 3.4.1.1.2) and the B8ZS line code.

Page 3-50

- R3-56** [58] The choice of AMI or B8ZS shall be provisionable by the user on a per-DS1 interface basis.

Page 3-50

- R3-57** [59] The asynchronous mapping of a DS1C into a VT3 SPE shall be as shown in Figure 3-30.

Page 3-51

- R3-58** [60] Twice in each VT3 SPE, the two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.

Page 3-51

- R3-59** [61] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single bit errors in the C-bits.
Page 3-51
- R3-60** [62] The stuffing mechanism that generates the C-bits shall be chosen so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 350 Hz, the output jitter is less than $1.0 UI_{pp}$ and $0.3 UI_{rms}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
Page 3-51
- R3-61** [904] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 350 Hz, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than the DS1C jitter transfer mask in Section 7.3.2 of GR-499-CORE.
Page 3-51
- R3-62** [63] The DS1C interface shall accommodate both the AMI line code (assuming the DS1C source meets the ones density criteria from GR-499-CORE of at least 12.5% ones over any 150 consecutive bits), and the B8ZS line code.
Page 3-51
- R3-63** [64] The choice of AMI or B8ZS shall be provisionable by the user on a per-DS1C interface basis.
Page 3-52
- R3-64** [65] The asynchronous mapping of a DS2 into a VT6 SPE shall be as shown in Figure 3-31.
Page 3-52
- R3-65** [66] Four times in each VT6 SPE, the two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.
Page 3-53
- R3-66** [67] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single-bit errors in the C-bits.
Page 3-53

R3-67 [68] The stuffing mechanism that generates the C-bits shall be chosen so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 500 Hz, the output jitter is less than $1.0 UI_{pp}$ and $0.3 UI_{rms}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.

Page 3-53

R3-68 [905] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 500 Hz, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than the DS2 jitter transfer mask in Section 7.3.2 of GR-499-CORE.

Page 3-53

R3-69 [69] The asynchronous mapping for a DS3 into an STS-1 SPE shall be as shown in Figure 3-32.

Page 3-55

R3-70 [70] In each subframe, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.

Page 3-55

R3-71 [71] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.

Page 3-55

R3-72 [906] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS3 jitter transfer mask shown in Figure 5-26, the output jitter is less than $0.4 UI_{pp}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.

Page 3-55

R3-73 [907] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS3 jitter transfer mask shown in Figure 5-26, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than that same DS3 jitter transfer mask.

Page 3-55

R3-74 [72] ATM cells shall be mapped into the STS-1 Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-1 SPE. The entire STS-1 Payload Capacity (i.e., 84 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 48.384 Mb/s.

Page 3-56

R3-75 [1044] The following apply if the HDLC-over-SONET mapping is supported.

- ... • The HDLC-framed signal shall be mapped into the STS-1 Payload Capacity by aligning the byte structure of every frame with the byte structure of the STS-1 SPE.
- ... • HDLC flags (i.e., '01111110' bytes) shall be used for interframe fill to account for the variable nature of the arrival of the HDLC frames.
- ... • The entire STS-1 Payload Capacity (i.e., 84 columns) shall be filled with HDLC frames and HDLC flags (as necessary).
- ... • The HDLC-framed signal plus the interframe fill shall be scrambled before it is inserted into the STS-1 Payload Capacity. In the reverse operation, after the STS-1 path is terminated the payload shall be descrambled before it is passed to the HDLC layer.
- ... • A self-synchronizing scrambler with a generator polynomial of $x^{43}+1$ shall be used.
- ... • The most significant bit of each byte of the HDLC-framed signal and interframe fill (i.e., the bit that will be placed into bit 1 of a byte in the STS-1 Payload Capacity, see Figure 3-2) shall enter the scrambler first, followed by the next most significant bit of that byte, etc.
- ... • The scrambler shall run continuously (e.g., it shall not be reset for each SONET or HDLC frame).

Page 3-57

O3-76 [1045] If the HDLC-over-SONET mapping is supported and a Cyclic Redundancy Check is applied over the HDLC payload signal, a CRC-32 should be used.

Page 3-58

R3-77 [73] The asynchronous mapping of a DS4NA into an STS-3c SPE shall be as shown in Figure 3-33.

Page 3-58

-
- R3-78** [74] In each row, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.
Page 3-59
- R3-79** [75] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.
Page 3-59
- R3-80** [76] The asynchronous mapping for a 125-Mb/s FDDI signal into an STS-3c SPE shall be as shown in Figure 3-34.
Page 3-61
- R3-81** [77] In each row, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.
Page 3-61
- R3-82** [78] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.
Page 3-61
- R3-83** [79] ATM cells shall be mapped into the STS-3c Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-3c SPE. The entire STS-3c Payload Capacity (i.e., 260 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 149.760 Mb/s.
Page 3-63
- R3-84** [80] DQDB slots shall be mapped into the STS-3c Payload Capacity by aligning the byte structure of every slot with the byte structure of the STS-3c SPE. The entire STS-3c Payload Capacity (i.e., 260 columns) shall be filled with slots, yielding a transfer capacity for DQDB slots of 149.760 Mb/s.
Page 3-63
- R3-85** [81] Bits 3 through 8 of the H4 byte shall contain a binary number in the range from '000000' (0) to '110100' (52) that indicates the offset between the H4 byte and the boundary of the first DQDB slot following the H4 byte.
Page 3-63
-

-
- R3-86** [82] Bits 1 and 2 of the H4 byte shall be used to carry the LSS.
Page 3-64
- R3-87** [83] The F2 and Z3 bytes of the STS POH shall carry the DQDB M1 and M2 bytes, respectively.
Page 3-64
- R3-88** [84] ATM cells shall be mapped into the STS-12c Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-12c SPE. The entire STS-12c Payload Capacity (i.e., 1040 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 599.040 Mb/s.
Page 3-65
- R3-89** [85] The pointer value shall be a binary number with a range of 0 to 782, and shall indicate the offset between the pointer word and the first byte of the STS SPE (as shown in Figure 3-38).
Page 3-67
- R3-90** [86] When there is a frequency offset between the frame rate of the Transport Overhead and that of the STS SPE, the pointer value shall be incremented or decremented as needed (although see rule ... 7. of **R3-100** [96v2]), accompanied by a corresponding positive or negative stuff byte.
Page 3-69
- R3-91** [87] A pointer increment operation shall be indicated by inverting bits 7, 9, 11, 13, and 15 (the I-bits) of the pointer word. The positive stuff byte shall appear immediately after the H3 byte in the frame containing the inverted I-bits, as shown in Figure 3-39.
Page 3-69
- R3-92** [88] A pointer decrement operation shall be indicated by inverting bits 8, 10, 12, 14, and 16 (the D-bits) of the pointer word. The H3 byte shall be used as the negative stuff byte, (i.e., it is used to carry an SPE byte in the frame containing the inverted D-bits), as shown in Figure 3-40.
Page 3-69
- R3-93** [89] If **O3-94** [90] (the "8 of 10" objective) is not met, then the increment decision shall be made by a majority vote of the I-bits, and the decrement decision shall be made by a majority vote of the D-bits.
Page 3-69
-

-
- O3-94** [90] The increment/decrement decision should be made at the receiver by a match of 8 or more of the 10 I- and D-bits to either the increment or decrement indication.
Page 3-69
- R3-95** [91] A normal NDF shall be indicated (during normal operation) by a '0110' code in the N-bits (see Figure 3-37). The NDF shall be set by inverting the N-bits to '1001.' The new alignment of the STS SPE shall be indicated by the pointer value accompanying the set NDF and takes effect at the offset indicated.
Page 3-72
- R3-96** [92] The decoding at the pointer processor shall be performed by majority voting (i.e., the NDF shall be detected as being set if three or four of the N-bits match the '1001' code). If a set NDF is detected, then the coincident pointer value shall replace the current value at the offset indicated by the new pointer value.
Page 3-72
- R3-97** [93] The first STS-1 within an STS-Nc shall have a normal pointer word.
Page 3-72
- R3-98** [94] All subsequent STS-1s within the STS-Nc shall have their pointer values (i.e., bits 7 through 16) set to all-ones, and their N-bits set to '1001' (i.e., set NDFs).
Page 3-72
- R3-99** [95] A pointer processor in an NE that is transmitting or receiving an STS-Nc SPE shall perform the operations indicated by the pointer in the first STS-1 of the STS-Nc on all N of the STS-1s in that STS-Nc.
Page 3-72
- R3-100** [96v2] The STS Payload Pointer shall be generated according to these rules:
- ... 1. During normal operation, a normal NDF is sent (i.e., the N-bits are set to '0110'), and the pointer value locates the start of the STS SPE within the STS Envelope Capacity.
 - ... 2. The pointer value shall only be changed by the operations in rules 4, 5, or 6.
 - ... 3. If an STS-Nc SPE is being transmitted, a normal pointer word is generated for the first STS-1 only. The concatenation indicator is
-

generated in the other pointers. All operations indicated by the pointer in the first STS-1 apply to each STS-1 in the STS-Nc.

- ... 4. If a positive stuff is needed, the current pointer value is sent with the I-bits inverted, and the subsequent positive stuff opportunity is considered an undefined byte. Subsequent pointers contain the previous pointer value incremented by one.
- ... 5. If a negative stuff is needed, the current pointer value is sent with the D-bits inverted, and the subsequent negative stuff opportunity is overwritten with an SPE byte. Subsequent pointers contain the previous pointer value decremented by one.
- ... 6. If the alignment of the SPE changes for any reason other than rules 4 or 5, the new pointer value shall be sent accompanied by a set NDF. The set NDF only appears in the first frame that contains the new value. The new SPE begins at the first occurrence of the offset indicated by the new pointer value.
- ... 7. No increment or decrement operation shall be performed for three frames following any of the operations in rules 4, 5, and 6.
- ... 8. For a nonterminated path, an incoming all-ones pointer word shall be regenerated or relayed with no more than a three-frame delay. When a non-all-ones pointer word is subsequently received, the downstream pointer shall be generated based on the pointer generation and interpretation criteria summarized in this requirement (R3-100 [96v2]) and R3-102 [97].

Page 3-73

- O3-101** [908] LTE that processes STS pointers should regenerate or relay an incoming all-ones pointer word with no more than a one-frame delay.

Page 3-74

- R3-102** [97] The STS Payload Pointer shall be interpreted according to these rules:

- ... 1. During normal operation, the pointer value locates the start of the STS SPE within the STS Envelope Capacity.
- ... 2. Any variation from the current pointer value shall be ignored unless a consistent new value is received three times consecutively, or the variation is one of the operations in rules 4, 5, or 6. Any consistent new value received three times in succession shall replace the current value at the offset indicated by the new pointer value.

- ...
3. If the pointer word contains the concatenation indicator, then the operations performed on that STS-1 are identical to those performed on the first STS-1 within the STS-Nc. Rules 4 and 5 do not apply to this pointer word.
- ...
4. If an increment is detected, then the byte following H3 shall be considered a positive stuff byte, and the current pointer value shall be incremented by one.
- ...
5. If a decrement is detected, then H3 shall be considered a negative stuff byte, and the current pointer value shall be decremented by one.
- ...
6. If a set NDF is detected, then the coincident pointer value replaces the current value at the offset indicated by the new pointer value.

Page 3-74

- R3-103** [98] The pointer value shall be a binary number with a range of 0 to 103 (VT1.5), 0 to 139 (VT2), 0 to 211 (VT3), or 0 to 427 (VT6), and shall indicate the offset between the pointer word and the first byte of the VT SPE (as shown in Figure 3-42).

Page 3-76

- R3-104** [99] When there is a frequency offset between the frame rate of the STS SPE and that of the VT SPE, the pointer value shall be incremented or decremented as needed (although see rule ... 6. of **R3-113 [108v2]**), accompanied by a corresponding positive or negative stuff byte.

Page 3-77

- R3-105** [100] A pointer increment operation shall be indicated by inverting bits 7, 9, 11, 13, and 15 (the I-bits) of the pointer word. The positive stuff byte shall appear immediately after the V3 byte in the superframe containing the inverted I-bits, as shown in Figure 3-42.

Page 3-77

- R3-106** [101] A pointer decrement operation shall be indicated by inverting bits 8, 10, 12, 14, and 16 (the D-bits) of the pointer word. The V3 byte shall be used as the negative stuff byte, (i.e., it is used to carry an SPE byte in the superframe containing the inverted D-bits), as shown in Figure 3-42.

Page 3-77

- R3-107** [102] If **O3-108 [103]** (the "8 of 10" objective) is not met, then the increment decision shall be made by a majority vote of the I-bits, and the decrement decision shall be made by a majority vote of the D-bits.

Page 3-77

- O3-108** [103] The increment/decrement decision should be made at the receiver by a match of 8 or more of the 10 I- and D-bits to either the increment or decrement indication.
Page 3-77
- R3-109** [104] Bits 5 and 6 of the VT Payload Pointer shall indicate the size of the VT using the code '00' (VT6), '01' (VT3), '10' (VT2), or '11' (VT1.5).
Page 3-77
- R3-110** [105] A normal NDF shall be indicated (during normal operation) by a '0110' code in the N-bits (see Figure 3-41). The NDF shall be set by inverting the N-bits to '1001.' The new alignment of the VT SPE shall be indicated by the pointer value accompanying the set NDF and takes effect at the offset indicated.
Page 3-78
- R3-111** [106] The decoding at the pointer processor shall be performed by majority voting (i.e., the NDF shall be detected as being set if three or four of the N-bits match the '1001' code). If a set NDF is detected, then the coincident pointer value shall replace the current value at the offset indicated by the new pointer value.
Page 3-78
- R3-112** [107] If a new size of VT is transmitted, then all 1 to 4 (depending on the new size) of the VT Payload Pointers in the VT group shall simultaneously indicate a set NDF and the same new size. The new size shall take effect immediately.
Page 3-78
- R3-113** [108v2] The VT Payload Pointer shall be generated according to these rules:
- ... 1. During normal operation, a normal NDF is sent (i.e., the N-bits are set to '0110'), the size bits indicate the size of the VT, and the pointer value locates the start of the VT SPE within the VT Envelope Capacity.
 - ... 2. The pointer value shall only be changed by the operations in rules 3, 4, or 5.
 - ... 3. If a positive stuff is needed, the current pointer value is sent with the I-bits inverted, and the subsequent positive stuff opportunity is considered an undefined byte. Subsequent pointers contain the previous pointer value incremented by one.

- ... 4. If a negative stuff is needed, the current pointer value is sent with the D-bits inverted, and the subsequent negative stuff opportunity is overwritten with an SPE byte. Subsequent pointers contain the previous pointer value decremented by one.
- ... 5. If the alignment of the SPE changes for any reason other than rules 3 or 4, the new pointer value shall be sent accompanied by a set NDF. The set NDF only appears in the first superframe that contains the new value. The new SPE begins at the first occurrence of the offset indicated by the new pointer value.
- ... 6. No increment or decrement operation shall be performed for three superframes following any of the operations in rules 3, 4, and 5.
- ... 7. For a nonterminated path, an incoming all-ones pointer word shall be regenerated or relayed with no more than a three-superframe delay. When a non-all-ones pointer word is subsequently received, the downstream pointer shall be generated based on the pointer generation and interpretation criteria summarized in this requirement (R3-113 [108v2]) and R3-115 [109].
- ... 8. If the size of the VTs within a VT group is to change, then the NDFs in all of the VTs of the new size (in that VT group) are set simultaneously.

Page 3-79

- O3-114** [909] STS PTE that processes VT pointers should regenerate or relay an incoming all-ones pointer word with no more than a one-superframe delay.

Page 3-80

- R3-115** [109] The VT Payload Pointer shall be interpreted according to these rules:

- ... 1. During normal operation, the pointer value locates the start of the VT SPE within the VT Envelope Capacity.
- ... 2. Any variation from the current pointer value shall be ignored unless a consistent new value is received three times consecutively, or the variation is one of the operations in rules 3, 4, or 5. Any consistent new value received three times in succession shall replace the current value at the offset indicated by the new pointer value.
- ... 3. If an increment is detected, then the byte following V3 shall be considered a positive stuff byte, and the current pointer value shall be incremented by one.

- ... 4. If a decrement is detected, then the V3 byte shall be considered a negative stuff byte, and the current pointer value shall be decremented by one.
- ... 5. If a set NDF is detected, then the coincident pointer value replaces the current value at the offset indicated by the new pointer value.
- ... 6. If the equipment has the capability to correctly process different VT sizes based on the received VT size bits, and a set NDF and an arbitrary new size of VT are received simultaneously in all of the VTs within a VT group, then the coincident pointer values and sizes shall replace the current pointer values and sizes at the offsets indicated in the new pointers.
- ... 7. If the equipment has the capability to correctly process different VT sizes based on the received VT size bits, then any variation from the current VT size shall be ignored unless consistent valid pointers indicative of a new VT size are received three times consecutively in all of the (new) VTs within a VT group, or the variation is the operation in rule 6. The VT size associated with such pointers received three times in succession shall replace the current size immediately.

Page 3-80

- R4-1** [110] For all SONET optical system interfaces described, binary Non-Return-to-Zero (NRZ) optical line coding shall be used.

Page 4-4

- R4-2** [111] To ensure the capability for upgrading SONET transport systems to high bit-rates, all splices (see GR-765-CORE, *Generic Requirements for Single Fiber Single-Mode Optical Splices and Splicing Systems*), connectors (see GR-326-CORE, *Generic Requirements for Single-Mode Optical Fiber Connectors*), attenuators (see GR-910-CORE, *Generic Requirements for Fiber Optic Attenuators*), couplers, and Wavelength-Division-Multiplexing (WDM) components (see GR-1209-CORE, *Generic Requirements for Fiber Optic Branching Components*) intended for installation in new facilities shall meet the reflectance requirements specified in the referenced documents.

Page 4-4

- R4-3** [112] The receiver reflectance shall be less than (more negative than) the value is listed under "Max. Receiver Reflectance" in Tables 4-3 through 4-11.

Page 4-5

- R4-4** [113] SONET optical transmitters and receivers shall operate properly in the presence of the worst-case combination of discrete reflectance, including receiver reflectance, system ORL, and minimum optical path attenuation values given in Tables 4-3 through 4-11. Proper system operation results in a system power penalty less than 1 dB under worst-case reflection conditions.
Page 4-5
- O4-5** [114] For all applications in Tables 4-3 through 4-11, the optical system should operate properly in the presence of a -8.5 dB reflection (i.e., maintain a system power penalty less than 1 dB).
Page 4-6
- R4-6** [115] The transmitter central wavelength and coupled transmit power shall be within the appropriate ranges listed in Tables 4-3 through 4-11.
Page 4-6
- R4-7** [116] The spectral width of the transmitter shall be less than or equal to the appropriate value listed in Tables 4-3 through 4-11.
Page 4-6
- R4-8** [117] The side-mode suppression ratio and extinction ratio of the transmitter shall be greater than or equal to the appropriate values listed in Tables 4-3 through 4-11.
Page 4-6
- O4-9** [118] In Tables 4-4, 4-5, and 4-9, it is an objective for MLM transmitters to meet the narrower spectral width specifications for those applications that list two possible values for $\Delta\lambda_{\text{rms}}$.
Page 4-6
- R4-10** [119] Transmit pulses, referenced to a noise filter with transfer characteristics as given in Equation 4-4, shall fall within the mask as defined in Figures 4-2 and 4-3.
Page 4-8
- R4-11** [120] The minimum acceptable value for receiver sensitivity shall equal the values P_{Rmin} specified in Tables 4-3 through 4-11.
Page 4-12
- R4-12** [121] The receiver shall accommodate an optical path power penalty of at least P_O for each application specified.
Page 4-12
-

-
- O4-13** [122] It is an objective that the receiver overload point equal or exceed the value of P_{Rmax} given in Tables 4-3 through 4-11. Page 4-12
- R4-14** [123] Suppliers shall provide worst-case values of transmission design parameters requested as part of system documentation. Page 4-23
- CR4-15** [124] A BCC may require that suppliers guarantee the worst-case values over the lifetime of their system components. Page 4-23
- R4-16** [125] The supplier shall provide general terminal equipment information in Worksheet 1. Page 4-23
- R4-17** [126] The supplier and the system integrator shall provide terminal equipment parameters under normal operating and short-term emergency conditions in Worksheets 2 and 3, respectively. Page 4-23
- R4-18** [127] If the terminal equipment has many options, the information of the worksheets shall be provided for each option. Page 4-23
- CR4-19** [128v2] A SONET NE may be required to support STS-1 electrical interfaces. Page 4-43
- CR4-20** [910] A SONET NE may be required to support STS-3 electrical interfaces. Page 4-43
- R4-21** [129] If a SONET NE supports STS-1 electrical interfaces, the following apply:
- ... • The transmitter shall generate an interface signal that meets the criteria in Table 4-12 for the entire range of interconnect cable lengths of 0 to 450 ft between the transmitter and the interface.
 - ... • The receiver shall accept any interface signal that conforms to the criteria in Table 4-12 (at the interface), and that propagates through a
-

jumper cable (if used, may be up to 27 ft) and an additional length of interconnect cable, which can also be up to 450 ft.

Page 4-43

R4-22 [130] If a SONET NE supports STS-3 electrical interfaces, the following apply:

- ... • The signal at the output of the transmitter (at the cable terminal) shall have a nominal rectangular pulse shape with a peak amplitude of 0.5 V (+10%) and maximum rise/fall times of 2 ns, as Figures 4-12 and 4-13 show.
- ... • The transmitter shall generate an interface signal that meets the criteria in Table 4-13 for the entire range of interconnect cable lengths of 0 to 225 ft between the transmitter and the interface.
- ... • The receiver shall accept any interface signal that conforms to the criteria in Table 4-13 (at the interface), and that propagates through a jumper cable (if used, may be up to 27 ft) and an additional length of interconnect cable, which can also be up to 225 ft.

Page 4-47

R5-1 [131] Before byte-interleaving to form an STS-N, the transport overhead byte positions of all the constituent STS-1s and STS-Ms shall be frame aligned.

Page 5-1

R5-2 [132] To form an STS-3 from STS-1s, three STS-1s shall be interleaved, one byte at a time. The first byte of the STS-3 shall be the A1 byte from the first STS-1, followed sequentially by the A1 byte from the second STS-1, and then the A1 byte from the third STS-1.

Page 5-2

R5-3 [133] To form a higher-level STS-N ($N > 3$) from lower-level STS-Ms ($3 \leq M < N$), the STS-Ms shall be interleaved $M/3$ bytes at a time. The output byte sequence shall be as shown in Figure 5-1.

Page 5-2

R5-4 [911] Each STS-Mc SPE in an STS-N shall be completely contained in one of Y groups of P STS-1s, where $Y = (N+P)$ and P is the smallest of the three numbers '3', '12', and '48' that satisfies the inequality $M \leq P \leq N$.

Page 5-2

-
- R5-5** [912] A SONET NE that provides the capability to terminate or pass a particular size STS-Mc SPE shall be capable of performing that function on an STS-Mc SPE that starts in any of the allowed starting positions defined in **R5-4** [911] and shown in Table 5-1.
Page 5-2
- R5-6** [134] SONET interface signals shall be scrambled (i.e., scrambled at the transmitter and descrambled at the receiver) using a frame synchronous scrambler of sequence length 127, operating at the line rate.
... The generating polynomial for the scrambler shall be $1+x^6+x^7$.
... The scrambler shall be reset to '111111' on the most-significant bit of the byte following the Z0 byte in the Nth STS-1 (i.e., the byte following the last Z0 byte). That bit and all subsequent bits to be scrambled shall be added, modulo 2, to the output from the x^7 position of the scrambler, as shown in Figure 5-3. The scrambler shall run continuously from that bit on throughout the remainder of the STS-N frame.
Page 5-6
- R5-7** [135] Overhead that is required to be generated shall carry valid data as this GR describes. Processing for the required overhead shall adhere to the criteria contained in this GR or the appropriate NE-specific GRs, TRs, and TAs.
Page 5-10
- R5-8** [1013] A SONET NE shall be capable of receiving and processing incoming SONET signals with bit-rates that are, at a minimum, anywhere in the range of ± 20 ppm off frequency from the nominal bit-rates for those signals.
Page 5-11
- CR5-9** [136] SONET NEs with line-side interfaces may be required to provide OW functionality.
Page 5-14
- O5-10** [137] If only a single OW channel is supported, it should be the LOW.
Page 5-14
- O5-11** [138] Access to the OW circuit should be through a 4-wire analog interface at 0 dBm. The input impedance should be $600\ \Omega$ ($\pm 5\%$), and the speech encoding should be μ -law Pulse Code Modulation (PCM).
Page 5-14
-

-
- | | |
|---------------|--|
| R5-12 | [139] If a 4-wire analog interface is not provided, either a 2-wire analog interface [0 dBm, 900 Ω ($\pm 5\%$), μ -law PCM], or a digital interface shall be provided. |
| | Page 5-14 |
| R5-13 | [140] The 8-bit PCM sample shall be synchronized to the STS-N frame, and the PCM bits shall be assigned to the corresponding bits of the appropriate E1 or E2 byte (see Figure 3-2). |
| | Page 5-15 |
| R5-14 | [141] The NE shall have the capability to generate the "quiet" PCM code (01111111) on its supported OW channels. |
| | Page 5-15 |
| R5-15 | [142] An NE with STE functionality but no LTE functionality shall provide the capability to pass the incoming LOW channel through to the outgoing LOW channel on the same line. |
| | Page 5-15 |
| CR5-16 | [143] A SONET ADM may be required to provide the capability to pass the EOW circuit between the high-speed OC-N signals. |
| | Page 5-15 |
| CR5-17 | [144] A SONET NE that terminates multiple SONET optical lines may be required to provide the capability to pass the EOW circuit between any two of those lines. |
| | Page 5-15 |
| CR5-18 | [145] A SONET NE may be required to support OW channel protection. |
| | Page 5-16 |
| R5-19 | [146] If OW channel protection is supported, then the OW protection scheme shall be the overhead protection scheme in which overhead channels are protected along with the traffic (see Section 8.3.1.3). |
| | Page 5-16 |
| CR5-20 | [147] An NE with STE functionality may be required to allow the user to access the Section User Channel (i.e., the F1 byte) in line-side signals. |
| | Page 5-16 |
| CR5-21 | [148] An NE that provides only STE functionality (e.g., an STE regenerator) may be required to be capable of passing the incoming F1 byte through to the outgoing signal on the same line. |
| | Page 5-16 |
-

-
- CR5-22** [149] An NE with STS PTE functionality may be required to allow the user to access the Path User Channel (i.e, the F2 byte).
Page 5-16
- CR5-23** [150] SONET LTE that terminates optical lines may be required to provide linear APS.
Page 5-17
- R5-24** [151] If linear APS is provided, the SONET NE shall provide the capability for the user to disable the feature on a per-interface basis.
Page 5-17
- R5-25** [152] Protection shall cover the multiplexer/optics units from the point, at or before, where the Line Overhead is inserted (the head end) to the point, at or beyond, where it is terminated (the tail end).
Page 5-17
- CR5-26** [153] The LTE may be required to support the 1+1 architecture.
Page 5-18
- R5-27** [154] If the 1+1 architecture is provided, the following apply:
- ... • The unidirectional mode shall be provided.
 - ... • If bidirectional switching is provided (see **CR5-28 [155]**), the mode shall be user-provisionable, with a default mode of unidirectional.
 - ... • If bidirectional switching is provided, the switching operation shall be bidirectional only if the LTE at both ends is provisioned to operate in the bidirectional mode. Otherwise, the LTE shall operate in the unidirectional mode. (The K2 byte is used to determine the mode of the far-end LTE as described in Section 5.3.5.2.)
 - ... • Nonrevertive switching shall be provided.
 - ... • If revertive switching is provided (see **CR5-29 [156]**), the choice of nonrevertive or revertive shall be user-provisionable, with a default of nonrevertive.
- Page 5-18
- CR5-28** [155] If the 1+1 architecture is provided, the bidirectional mode may be required to be provided.
Page 5-19
-

- CR5-29** [156] If the 1+1 architecture is provided, revertive switching may be required to be provided.
Page 5-19
- CR5-30** [157] 1+1 LTE may be required to be upgradable to the 1:n architecture.
Page 5-19
- CR5-31** [158] LTE may be required to support the 1:n architecture.
Page 5-19
- CR5-32** [159] LTE supporting the 1:n architecture may be required to support values of n greater than 1.
Page 5-19
- R5-33** [160] If the 1:n architecture is provided, the following apply:
- ... • All switching shall be revertive.
 - ... • The unidirectional mode shall be provided.
 - ... • The bidirectional mode shall be provided.
 - ... • The mode shall be user-provisionable, with a default mode of bidirectional.
 - ... • The switching operation shall be unidirectional only if the LTE at both ends is provisioned to operate in the unidirectional mode. Otherwise, the LTE shall operate in the bidirectional mode. (The K2 byte is used to determine the provisioned mode of the far-end LTE as described in Section 5.3.5.2.)
- Page 5-19
- O5-34** [161] If the 1:n architecture is provided, the LTE should provide the capability to transport extra traffic on the protection line when it is not being used for protection.
Page 5-20
- R5-35** [162] If the capability to transport extra traffic is provided, the SONET NE shall provide the capability for the user to disable that feature for interworking with other SONET NEs that do not support it.
Page 5-20
- CR5-36** [163] LTE that supports only the 1:1 case of the 1:n architecture may be required to be upgradable to support values of $n > 1$.
Page 5-20
-

-
- R5-37** [164] 1:1 LTE (which indicates the 1:n architecture on bit 5 of the transmitted K2 byte) shall operate as 1+1 LTE if the far-end LTE indicates that it is 1+1 LTE, as detected on the received K2 byte. The 1:1 LTE shall continue to indicate the 1:n architecture on the transmitted K2 byte unless it is reprovisioned by the user to 1+1. It shall also continue to indicate its provisioned unidirectional/bidirectional switching mode; however, it shall meet the criteria in Section 5.3.2.1 (instead of the criteria in Section 5.3.2.2) concerning which of those modes to actually operate in.
Page 5-20
- R5-38** [165] Loss of Signal, Loss of Frame and AIS-L defects (see Section 6.2.1), and a Line BER exceeding 10^{-3} on an incoming OC-N shall be detected as SF conditions on that line.
Page 5-21
- CR5-39** [166] The BER threshold for an SF condition may be required to be user-provisionable over the range of 10^{-3} to 10^{-5} .
Page 5-21
- R5-40** [167] A BER exceeding the SD threshold on an incoming OC-N shall be detected as an SD condition on that line.
Page 5-21
- R5-41** [168] The BER threshold for an SD condition shall be user-provisionable over the range of 10^{-5} to 10^{-9} .
Page 5-21
- R5-42** [169] For SF conditions caused by LOS, LOF, or AIS-L defects, the switch initiation time shall be 10 ms or less.
Page 5-21
- O5-43** [170] For SF conditions caused by LOS, LOF, or AIS-L defects, the switch initiation time should be 8 ms or less.
Page 5-21
- R5-44** [171] For SF and SD conditions based on BER, the switch initiation time shall be below the "maximum" curve in Figure 5-5 (assuming the actual BER is greater-than or equal to the threshold).
Page 5-22
- O5-45** [172] For SF and SD conditions based on BER, the probability that the switch initiation time will be less than the "objective" curve in Figure 5-5
-

(for the particular rate OC-N signal) should be greater than 0.95 (assuming the actual BER is greater than or equal to the threshold).

Page 5-22

- R5-46** [173v2] For an SF or SD detection threshold of 10^{-n} and an actual BER of $1 \times 10^{-(n+1)}$ or less, the probability that the SD or SF condition will be detected in the "maximum" switch initiation time from Figure 5-5 (for that particular threshold) shall be less than or equal to 10^{-6} .

Page 5-22

- R5-47** [174] The time to complete a switch, once it is initiated, shall be 50 ms or less.

Page 5-23

- R5-48** [175] For bidirectional switching, the LTE at both ends shall complete the switch within the same 50-ms switch completion time, from the time the request is initiated.

Page 5-23

- CR5-49** [913] Manually initiated facility protection switches may be required to be error-free.

Page 5-24

- R5-50** [176] The clearing threshold for an SD or SF condition based on the BER shall be one-tenth the threshold for declaring the SD or SF.

Page 5-25

- R5-51** [1046] If an SD or SF condition has been detected and the incoming signal's BER is greater than or equal to that SD or SF threshold, the probability that the LTE will detect that the BER is less than the SD or SF clearing threshold within the "maximum clearing" time listed in Table 5-3 shall be less than or equal to 10^{-6} .

Page 5-25

- R5-52** [1047] If an SD or SF condition has been detected and the incoming signal's BER is less than or equal to the SD or SF clearing threshold, the probability that the LTE will detect that the BER is less than that threshold within the "maximum clearing" time shown in Table 5-3 shall be greater than or equal to 0.99.

Page 5-25

- O5-53** [177v3] If an SD or SF condition has been detected and the incoming signal's BER is less than or equal to the SD or SF clearing threshold, the probability that the LTE will detect that the BER is less than that threshold

within the "objective clearing" time shown in Table 5-3 should be greater than 0.95.

Page 5-25

- R5-54** [914] Once LTE has detected that the BER is less than the SD or SF clearing threshold, it shall clear the SD or SF condition within an additional 10.5 seconds (although see below for possible exceptions in cases of intermittent SD and SF conditions) assuming the LTE does not detect that the BER is greater than or equal to the SD or SF threshold before that condition is cleared.

Page 5-26

- CR5-55** [915] LTE may be required to be designed to reduce the chance or number of rapid protection switching oscillations that could occur in multiple failure or degradation situations where one or more of the failures or degradations is intermittent.

Page 5-27

- R5-56** [916] If LTE is designed to reduce the chance or number of rapid protection switching oscillations, the method used shall be clearly documented.

Page 5-27

- R5-57** [178v2] For LTE using revertive switching, a WTR period of 5 to 12 minutes shall be provided after the condition that caused an automatically initiated switch to the protection line clears. The length of the WTR period shall be user-provisionable on a per-protection line (or per-protection group) basis.

Page 5-28

- R5-58** [179] Bits 1 through 4 of the K1 byte shall indicate the current request using the codes listed in Table 5-4.

Page 5-29

- R5-59** [180] An NE using the 1:n architecture shall provide the capability to provision each working channel and the null channel (for conditions detected on the Protection line) as high or low priority, with low priority as the default. These priorities shall determine which of the listed codes is used for SF and SD requests.

Page 5-29

- R5-60** [181] Bits 5 through 8 of the K1 byte shall indicate the number of the channel for which the request is issued, using the codes shown in Table 5-5.

Page 5-31

- R5-61** [182] All local requests [i.e., any locally detected SF or SD conditions, local WTR, Do Not Revert (DNR) or No Request state, or external request from a received switch command] shall be evaluated to determine the highest priority local request based on the order of request priorities in Table 5-4. If local SF or SD conditions of the same priority have been detected and are still present on different lines at the same time, then the condition with the lowest channel number shall take precedence in the evaluation.

... The highest priority local request shall be compared to the current local request. If the highest priority local request is of higher priority than the current local request (based on the order of priorities in Table 5-4) or if the current local request is no longer a valid request (e.g., the condition or external request that caused it has been cleared) then the highest priority local request shall replace the current local request (i.e., it becomes the new current local request). In all other cases the current local request shall not change.

Page 5-31

- R5-62** [183] In the bidirectional mode, the priorities of the current local request and the remote request on the received K1 byte shall be compared according to the order of priorities in Table 5-4. A received Reverse Request shall not be considered in the comparison, because it assumes the priority of the request to which it is responding. The transmitted K1 byte shall be set to indicate a Reverse Request if any of the following are true:

- ... • The remote request is of higher priority
- ... • The requests are of the same priority, they are higher priority than a No Request, and the transmitted K1 byte is already set to indicate Reverse Request
- ... • The requests are of the same priority, they are higher priority than a No Request, the transmitted K1 byte is not set to indicate Reverse Request, and the remote request indicates a lower channel number.
- ... The transmitted K1 byte shall be set to indicate the local request in all other cases.

Page 5-32

-
- R5-63** [184] In the unidirectional mode, the transmitted K1 byte shall be set to indicate the current local request (i.e., Reverse Request is never indicated).
Page 5-33
- R5-64** [185] For working channels at LTE using revertive switching, when a local condition that caused an automatically initiated switch clears, a local Wait-to-Restore (WTR) state shall be activated.
Page 5-33
- R5-65** [186] A WTR state shall normally time out and become a No Request – null channel (or No Request – Channel 15, if applicable). The WTR timer shall deactivate earlier if the transmitted K1 byte no longer indicates WTR (i.e., when any request of higher priority preempts this state). When the higher priority request is cleared, the preempted WTR state shall not be reactivated. (Note however, that a new WTR state would be required to be activated if the higher priority request was for an automatically initiated switch of a working channel.)
Page 5-33
- R5-66** [187] When an external request is cleared, the No Request – null channel (or No Request – Channel 15, if applicable) state shall be activated (i.e., the WTR state is not activated).
Page 5-33
- R5-67** [188] For the working channel at LTE using nonrevertive switching, the selection of the working channel from the protection line shall be maintained by activating a Do Not Revert (DNR) state (instead of a WTR or No Request state). The DNR state shall be deactivated if the transmitted K1 byte no longer indicates DNR (i.e., when any request of higher priority preempts this state).
Page 5-33
- R5-68** [189] After any request for the null channel is cleared, the No Request – null channel (or No Request – Channel 15, if applicable) state shall be activated.
Page 5-34
- R5-69** [190v2] The bit assignments for the K2 byte shall be as follows (see Section 5.3.5.2.2 for the corresponding K2 byte generation rules):
- ... • Bits 1 through 4 – A channel number using the codes shown in Table 5-5.
 - ... • Bit 5 – Indication of architecture (1+1 or 1:n), as provisioned.
-

- ... • Bits 6 through 8 – Indication of the mode of operation, as provisioned, or non-APS channel uses (i.e., AIS-L, RDI-L).

Page 5-34

R5-70 [191] For all architectures and modes of operation, bits 1 through 4 of the K2 byte shall be set to indicate:

- ... • The null channel (0) if the received K1 byte indicates the null channel
- ... • The number of the channel bridged onto the protection line in all other cases.

Page 5-34

R5-71 [192] Bit 5 of the K2 byte shall be set to indicate:

- ... • Code '0' if the provisioned (or only supported) architecture is 1+1
- ... • Code '1' if the provisioned (or only supported) architecture is 1:n.

Page 5-35

R5-72 [193] Bits 6 through 8 of the K2 byte shall be set to indicate:

- ... • Code '101' if the provisioned mode is bidirectional
- ... • Code '100' if the provisioned (or only supported) mode is unidirectional.

Page 5-35

R5-73 [194] An APS Mode Mismatch indication resulting from a mismatch of K2 bit 5 shall be used to modify the operation of the 1:1 LTE to interwork with the 1+1 LTE (see Section 5.3.2.3).

Page 5-36

R5-74 [195] An APS Mode Mismatch indication resulting from a mismatch of K2 bits 6 through 8 shall be used by 1+1 LTE provisioned for bidirectional switching to operate unidirectionally, or by 1:n LTE provisioned for unidirectional switching to operate bidirectionally (see Sections 5.3.2.1 and 5.3.2.2).

Page 5-36

R5-75 [196] If LTE stops receiving an indication of the provisioned mode of operation from the far-end LTE, then it shall maintain its current mode of operation.

Page 5-36

- R5-76** [197v2] In the 1:n architecture, the channel whose number is indicated on the received K1 byte shall be bridged to the protection line unless the request is invalid (e.g., in the bidirectional mode, the requesting channel is locked out locally).
Page 5-36
- R5-77** [198] If a local SF condition is detected on the protection line or a Protection Switching Byte failure is declared, the current bridge shall be maintained if the mode of operation is unidirectional, or shall be released if the mode is bidirectional.
Page 5-36
- R5-78** [199] In the 1+1 architecture, the working channel shall be continuously bridged to the protection line.
Page 5-36
- R5-79** [200] In all architectures and modes except the 1+1 unidirectional mode, if there is a match of the transmitted K1 and received K2 bytes, then the indicated channel shall be selected from the protection line, unless one of the following is true (in which case the selector shall be in the released position, see Figure 5-6):
- ... • The match is for the null channel
 - ... • An Exercise request is indicated on the transmitted K1 byte (unidirectional and bidirectional), or the received and acknowledged K1 byte (bidirectional only).
- Page 5-38
- R5-80** [201] In the 1:n architecture, the selector shall also be in the released position when there is a mismatch of the channel numbers.
Page 5-38
- R5-81** [202] In the 1+1 unidirectional mode, the working channel shall be selected from the protection line if channel number 1 is indicated on the transmitted K1 byte.
Page 5-38
- R5-82** [203] The linear APS protocol shall be carried between LTE in the APS channel (i.e., the K1 and K2 bytes) transmitted on the protection line.
Page 5-39

-
- R5-83** **[204]** A new code on the received K1 and K2 bytes shall replace the current received code if it is received identically in three consecutive frames.
Page 5-39
- R5-84** **[205]** LTE shall not transmit invalid codes.
Page 5-39
- R5-85** **[206]** If the capability to transport extra traffic on the protection line is provided, the No Request code shall be used to keep the extra traffic channel on the protection line (i.e., the No Request code is the only valid code to transmit with a channel number of '1111').
Page 5-39
- O5-86** **[917]** LTE should not consider an SD or SF request detected on the incoming K1 bits 1 through 4 to be invalid based (solely) on the high/low priority of that request.
Page 5-40
- R5-87** **[207]** LTE operating in the 1+1 bidirectional mode and using nonrevertive switching shall consider the WTR code for the working channel to be valid.
Page 5-40
- O5-88** **[208]** LTE operating in the 1+1 bidirectional mode should consider the Do Not Revert code for either the null channel or the working channel as valid.
Page 5-40
- R5-89** **[209]** Even when accepted as the current code, an invalid code in K1 shall not result in any immediate protection switching action.
Page 5-41
- R5-90** **[210]** For LTE operating in the bidirectional mode, the protection line shall be considered to be in the SF condition when a Protection Switching Byte failure is declared. An SF condition resulting from a Protection Switching Byte failure shall be cleared when the Protection Switching Byte failure is cleared.
Page 5-41
- R5-91** **[211]** The following switch commands shall be provided, as described.
... **Clear** – Clears all of the switch commands listed below, for the channel or channels specified in the command.
-

- ... **Lockout of Protection** – Prevents any of the working channels from switching to the protection line by issuing a Lockout of Protection request [unless a request of equal priority (i.e., a Lockout of Protection) is already in effect].
- ... **Forced Switch of Working (to Protection)** – Switches the specified working channel to the protection line unless a request of equal or higher priority is in effect by issuing a Forced Switch request.
- ... **Forced Switch of Protection (to Working)** – Switches the working channel back from the protection line to the working line unless a request of equal or higher priority is in effect, by issuing a Forced Switch request for the null channel. This command applies only in the 1+1 architecture.
- ... **Manual Switch of Working (to Protection)** – Switches the working channel to the protection line unless a request of equal or higher priority is in effect, by issuing a Manual Switch request.
- ... **Manual Switch of Protection (to Working)** – Switches the working channel back from the protection line to the working line unless a request of equal or higher priority is in effect, by issuing a Manual Switch request for the null channel. This command applies only in the 1+1 architecture.

Page 5-42

- R5-92** [212] When a higher priority local or remote request preempts an external request, the preempted request shall not be retained (i.e., when the higher priority request is cleared, the preempted switch request shall not be reinitiated).

Page 5-42

- CR5-93** [213] LTE capable of operating in the 1+1 bidirectional mode or the 1:n architecture may be required to support the Exercise command.

Page 5-43

- R5-94** [214] If the Exercise command is supported, it shall cause the LTE to perform as described below.

- ... **Exercise** – Exercises the protocol for a protection switch of the specified channel, unless a request of equal or higher priority is in effect, by issuing an Exercise request for that channel and checking the response on the APS channel.

Page 5-43

- R5-95** [215] If the Exercise command is supported, it shall be cleared automatically at the end of the Exercise routine or the required switch completion time, whichever is sooner.
Page 5-43
- R5-96** [216] LTE that does not support the Exercise command shall consider an incoming Exercise request with a valid channel number to be valid, and shall respond as requested (per Section 5.3.5).
Page 5-43
- R5-97** [217v2] The following control commands shall be supported by 1:n LTE, and shall cause the LTE to perform as described below.
- ... **Lockout a Working Channel** – Prevents the specified working channel (or channels) from switching to the protection line.
- ... **Clear Lockout-a-Working-Channel** – Clears the Lockout a Working Channel command for the channel or channels specified in the clear command.
Page 5-43
- R5-98** [918] The physical interface between the SONET NE and the synchronization network shall meet the criteria in Section 3.2.2 of GR-1244-CORE.
Page 5-52
- R5-99** [224v2] SONET LTE shall generate and provide the capability to process the synchronization status messages listed in Table 5-7 on bits 5 through 8 of the S1 bytes of all signals at SONET line-side “interfaces” (except for lines 2 through n at an OC-N “interface” where 1:n linear APS is being used).
Page 5-54
- R5-100** [225v3] SONET LTE shall generate the synchronization status messages listed in Table 5-7 on bits 5 through 8 of the S1 bytes of all signals at SONET drop-side “interfaces”.
Page 5-54
- CR5-101** [1048] SONET LTE may be required to provide the capability to process the synchronization status messages listed in Table 5-7 on bits 5 through 8 of the S1 bytes of all signals at SONET drop-side “interfaces”.
Page 5-54

-
- CR5-102** [1049] A SONET NE that contains LTE and supports line-timing (or through-timing), or that is capable of being provisioned to derive DS1s from the incoming signals at one or more of its SONET "interfaces", may be required to be able to be provisioned by the user to ignore the incoming S1 byte at its provisioned reference and/or derived DS1 source "interfaces".
Page 5-54
- R5-103** [1050] If a SONET NE allows the user to disable the processing of the S1 byte at its SONET "interfaces" that are provisioned as timing references or derived DS1 sources, the default shall be that the S1 byte is processed.
Page 5-54
- R5-104** [222] The "Reserved for Network Synchronization" message shall be treated as a "DON'T USE for Synchronization" message at intercarrier interfaces.
Page 5-56
- R5-105** [223] Any proprietary use of the "Reserved for Network Synchronization Use" message shall be clearly documented as per Section 3.2.
Page 5-56
- R5-106** [226] For NEs that support the external-timing mode, that mode shall be the default timing mode.
Page 5-56
- R5-107** [227] For NEs that support automatic switching between timing modes, the switching shall be revertive.
Page 5-57
- R5-108** [919] SONET NEs with external-timing interfaces shall meet the criteria in Section 3.2.1 of GR-1244-CORE.
Page 5-57
- R5-109** [920] SONET NEs with line-timing "interfaces" shall meet the criteria in Section 3.2.3 of GR-1244-CORE.
Page 5-58
- R5-110** [1051] If an NE that supports line APS does not support provisioning of an OC-N/M "interface" as a single reference, that fact shall be clearly documented.
Page 5-58
-

CR5-111 [239] In some applications the NE may be required to provide the user the capability to provision a line-side OC-M "interface" as a synchronization source.

Page 5-58

R5-112 [240v2] When an NE is through-timed, the transmitted signals at the "west" OC-N or STS-N electrical "interface" shall be timed by the terminating signals at the "east" OC-N or STS-N electrical "interface", and the transmitted signals at the "east" "interface" shall be timed by the terminating signals at the "west" "interface".

Page 5-60

R5-113 [241] An ADM that supports through-timing shall provide the user with the capability to provision for automatic switching to line-timing using the non-failed OC-N "interface" when one of the OC-N "interfaces" becomes unavailable as a reference, as defined in Section 5.4.6.

Page 5-60

R5-114 [242] If a through-timed ADM has OC-M, STS-M electrical, or synchronous DS1 interfaces, the timing source for these outputs, as a group or individually, shall be user-provisionable from either of the OC-N "interfaces".

Page 5-60

CR5-115 [243v2] Some service providers may require stratum 3 clocks for SONET NEs used in applications that do not explicitly require stratum 3 clocks.

Page 5-61

CR5-116 [244v2] Some providers may require NEs to provide stratum 3E clocks in certain applications, such as NEs that serve as timing distribution hubs.

Page 5-61

CR5-117 [245v2] Some providers may require NEs to provide stratum 2 clocks in certain applications, such as NEs that serve as timing distribution hubs.

Page 5-61

R5-118 [246] A stratum 3, 3E or 2 clock in a SONET NE shall meet the following criteria in GR-1244-CORE:

- ... • minimum free-run accuracy (GR-1244-CORE, Section 5.1)
- ... • holdover stability (GR-1244-CORE, Section 5.2)
- ... • pull-in/hold-in range (GR-1244-CORE, Section 3.5)

Page 5-61

-
- R5-119** [247] A stratum 3E or 2 clock in a SONET NE shall meet the following requirements in GR-1244-CORE:
- ... • wander tolerance (GR-1244-CORE, Section 4.3)
 - ... • wander transfer (GR-1244-CORE, Section 5.4)
- Page 5-61
- R5-120** [921] A stratum 3 clock in a SONET NE shall meet the SMC wander transfer requirements in Section 5.4.4.2.4 (of this document).
- Page 5-61
- R5-121** [248] The minimum free-run accuracy of an SMC shall be ± 20 ppm.
- Page 5-62
- R5-122** [249v3] A SONET NE containing an SMC shall be capable of entering holdover when all of its timing references are determined to be failed (as per Section 5.4.6) or contain the "DON'T USE for Synchronization" synchronization status message.
- Page 5-62
- R5-123** [922] Any transient associated with entry into holdover shall be bounded by the MTIE mask in Figure 5-14.
- Page 5-62
- R5-124** [923] The initial fractional frequency offset, as defined in T1.105.09, shall be less than 0.05 ppm.
- Page 5-62
- R5-125** [924] The frequency drift rate, as defined in T1.105.09, shall be less than 5.8×10^{-6} ppm/second.
- Page 5-62
- R5-126** [925] The fractional frequency offset under varying temperature conditions shall not exceed 4.1 ppm.
- Page 5-62
- R5-127** [250] Entry into holdover, and restoration from holdover, shall be error-free.
- Page 5-63
-

-
- R5-128** [253] If a SONET NE with an SMC is timed from an external reference, the NE clock shall pull-in and hold-in to an external reference that is off frequency by ± 4.6 ppm (i.e., from a free-running stratum 3 clock).
Page 5-64
- R5-129** [254] If a SONET NE with an SMC is timed from an OC-N reference, the NE clock shall pull-in and hold-in to an OC-N that is off frequency by ± 20 ppm.
Page 5-64
- R5-130** [926] The maximum settling time for an SMC shall be 100 seconds.
Page 5-64
- R5-131** [255v2] OC-N/OC-M and STS-N/STS-M electrical outputs, when referenced to an external timing signal that meets the wander TDEV mask in Figure 5-16, shall be "less than or equal to" (where "less than or equal to" is defined to allow a phase gain of up to 0.2 dB in the pass-band of the clock) the wander TDEV mask given in Figure 5-15.
Page 5-64
- R5-132** [256v2] OC-N/OC-M and STS-N/STS-M electrical outputs, when referenced to an OC-N timing signal that meets the wander TDEV mask in Figure 5-15, shall be "less than or equal to" (where "less than or equal to" is defined to allow a phase gain of up to 0.2 dB in the pass-band of the clock) the wander TDEV mask in Figure 5-15.
Page 5-65
- R5-133** [264v2] A stratum or SMC clock in a SONET NE shall meet the duplex equipment criteria in Section 3.3 of GR-1244-CORE.
Page 5-67
- R5-134** [257] OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the MTIE wander mask in Figure 5-17 when timed with a wander-free reference.
Page 5-67
- R5-135** [258] OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the TDEV wander mask in Figure 5-18 when timed with a wander-free reference.
Page 5-67
- R5-136** [259v2] For all SONET NEs that contain LTE, OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the specifications in ANSI T1.101-1994 for OC-N phase transients during synchronization
-

rearrangement operations. Those specifications specify an MTIE of no greater than the "requirement" mask in Figure 5-19. Rearrangement activities include the following:

- ... • Manual timing reference switching
- ... • Automatic timing reference switching as described in Section 5.4.6
- ... • Switching between working line 1 and the protection line at the active reference OC-N "interface" (for NEs that support line APS)
- ... • Entry into self-timing operation (i.e., holdover or free-run) for the initial 2.33 seconds of self-timing
- ... • Automatic clock diagnostics
- ... • Clock hardware protection switching.
- ... • Phase transients on an external or OC-N synchronization input with the rate of change as specified in ANSI T1.101-1994

Page 5-70

- R5-137** [1014] For SONET NEs that contain stratum 2 internal clocks, the MTIE of the SONET outputs during the internal rearrangements listed above (i.e., the first six rearrangements listed) shall meet the "objective" mask in Figure 5-19.

Page 5-70

- O5-138** [260] For SONET NEs that contain stratum 3E, stratum 3 or SMC internal clocks, the MTIE of the SONET outputs during phase transients caused by the internal rearrangements listed above should be no greater than the "objective" mask in Figure 5-19.

Page 5-70

- O5-139** [928] The SONET outputs of an NE should meet the jitter generation requirement in Section 5.6.2.3.6 during the synchronization rearrangement activities listed in Section 5.4.4.3.3.

Page 5-72

- R5-140** [261v2] Except for clock hardware protection switching, the synchronization rearrangement activities listed in Section 5.4.4.3.3 shall cause no errors on payload traffic.

Page 5-72

- O5-141** [262v2] Clock hardware protection switching should cause no errors on payload traffic.

Page 5-72

-
- R5-142** [265] Recovery from self-timing (holdover or free-run) shall be automatic.
Page 5-72
- CR5-143** [266] An NE with a stratum clock or SMC may be required to provide the ability to inhibit automatic recovery from self-timing.
Page 5-72
- R5-144** [267v2] Automatic restoration from holdover shall conform to the criteria in GR-1244-CORE, Sections 3.6 and 3.7. An SMC shall conform with the criteria for a stratum 3 clock.
Page 5-72
- R5-145** [268] Automatic restoration from the free-run mode shall occur within two seconds of the presence of a validated reference signal.
Page 5-72
- R5-146** [930v2] The maximum rate of frequency change during holdover recovery shall be less than 2.9 ppm/second.
Page 5-73
- R5-147** [271v2] For interruptions (i.e., short LOS, AIS, OOF or LOF defects, not phase or frequency transients) of reference signals that do not cause reference switches or switches between lines (at an NE that supports line APS), the output criteria of Figure 5-17 shall be met.
Page 5-73
- R5-148** [272] The NE shall tolerate phase transients on external and OC-N reference signals of a magnitude and slope as defined in ANSI T1.101-1994.
Page 5-73
- R5-149** [1015] Clocks synchronized to an external DS1 timing signal shall tolerate, as a minimum, the jitter specified for the input test signal in the wander generation requirements in Section 5.4.4.3.2.
Page 5-73
- CR5-150** [1016] Clocks synchronized to an external DS1 timing signal may be required to tolerate, as a minimum, input jitter applied according to the mask in Figure 7-1 of GR-499-CORE.
Page 5-73
-

- R5-151** [1017] Clocks that are lined-timed from an incoming OC-N signal shall meet the Category II jitter tolerance requirements in Figure 5-28.
Page 5-73
- R5-152** [273v3] A SONET NE that contains LTE shall have the capability to supply two DS1 timing reference signals. The NE shall be capable of deriving both of these DS1s from a single line-side OC-N "interface" (see Figure 5-20) and, if more than one OC-N "interface" is supported, of deriving each DS1 from a different OC-N "interface" (see Figure 5-21). As a minimum, the derived DS1 signals shall be in the Superframe format and shall meet the ones-density requirement in GR-499-CORE.
Page 5-74
- CR5-153** [274v2] An NE that supports line APS may be required to switch the source of timing for a derived DS1 to the OC-N on the protection line when the OC-N on working line 1 becomes unavailable as a derived DS1 source due to an LOS, LOF, or AIS (and vice versa if one of the DS1s is normally derived from the OC-N on the protection line).
Page 5-76
- CR5-154** [1018] An NE that provides more than one OC-N "interface" may be required to support switching of the source of timing for a derived DS1 to a different (secondary) OC-N "interface" when the OC-N signal(s) at the original (primary) "interface" becomes unavailable as a derived DS1 source due to an LOS, LOF, or AIS, or when switching is appropriate based on the received synchronization status messages (see **R5-169** [285v3]).
Page 5-76
- R5-155** [1019] If an NE that provides line APS supports switching between OC-N "interfaces" as the source of timing for a derived DS1, then it shall conform to **CR5-153** [274v2]. In addition, a switch between "interfaces" in response to an LOS, LOF, or AIS shall occur only if the OC-N signals on both working line 1 and the protection line have failed.
Page 5-76
- R5-156** [1020] A line-timed NE that is provisioned to support switching between OC-N "interfaces" as the source of timing for a derived DS1, and to (as a default) derive all of its active DS1s from the same OC-N "interface", shall use the same type of switching (i.e., nonrevertive or revertive) as it uses for timing reference switching.
Page 5-76

R5-157 [1021] A through-timed ADM that is provisioned to support switching between OC-N "interfaces" as the source of timing for a derived DS1, and to (as a default) derive each DS1 from a different OC-N "interface", shall use revertive switching.

Page 5-77

CR5-158 [1052] A SONET NE that supports the capability to switch between OC-N "interfaces" as the source for a derived DS1 may be required to support derived DS1 source switching-related commands equivalent to the timing reference switching-related commands (defined in Section 5.4.6) that it supports.

Page 5-77

R5-159 [1053] If a SONET NE supports one or more derived DS1 source switching-related commands, the effects of those commands shall be functionally equivalent to the effects of the corresponding timing reference switching-related commands.

Page 5-77

O5-160 [276] The derived DS1 should be a framed all-ones signal.

Page 5-78

CR5-161 [277] The NE may be required to provide the capability for the user to provision the derived DS1 in the ESF format (in addition to the required Superframe format).

Page 5-78

R5-162 [278] The SONET NE shall be capable of supporting all its timing modes when providing derived DS1s from an OC-N.

Page 5-78

R5-163 [279v2] DS1 AIS (unframed all-ones; not the A-bit code for ESF) shall be inserted into the derived DS1 when the OC-N signal becomes unavailable as a derived DS1 source due to an LOS, LOF, or AIS (assuming no switching, see **CR5-153** [274v2], **CR5-154** [1018], **R5-155** [1019] and **R5-169** [285v3]). DS1 AIS shall be generated no later than the declaration of failure (see Section 6.2.1).

Page 5-78

R5-164 [280v2] Automatic restoration of the derived DS1 shall occur within two seconds of the clearing of the derived DS1 source failure.

Page 5-78

-
- R5-165** [281] The MTIE of the derived DS1 shall be less than 50 ns for observation times beyond the jitter region (observation times longer than 0.1 second).
Page 5-78
- R5-166** [282] The TDEV of the derived DS1 shall be below the mask in Figure 5-22.
Page 5-78
- R5-167** [283] The jitter on the derived DS1 shall be less than $1.0 U_{I_{pp}}$.
Page 5-79
- R5-168** [284v2] The derived DS1 shall meet the MTIE requirement for DS1 phase transients in ANSI T1.101 during rearrangements (e.g., switching the derived DS1 source between working line 1 and the protection line in a system that supports line APS, or between "east" and "west" OC-Ns as per Section 5.4.5.2.1). For observation periods up to 280 seconds, a phase transient on a derived DS1 shall not exceed a magnitude of 1 μ s with a slope of 81 ns for a measurement period of 1.326 ms.
Page 5-79
- R5-169** [285v3] A line-timed NE or through-timed ADM shall automatically select (from the OC-N "interfaces" provisioned as possible sources for each derived DS1) the OC-N "interface" with the highest quality synchronization status message as the source for that derived DS1 signal.
Page 5-81
- R5-170** [286] An NE shall support the "threshold AIS generation" mode.
Page 5-82
- CR5-171** [287v2] An NE may be required by some service providers to support the "message pass-through" mode (i.e., it may be required to support synchronization status messages on the derived DS1).
Page 5-82
- R5-172** [288] If an NE supports both message translation modes, the format of the derived DS1 signal shall indicate which mode is used. If the derived DS1 is in the ESF format, the "message pass-through" mode shall be used. If the derived DS1 is in the SF format, the "threshold AIS generation" mode shall be used.
Page 5-82
- R5-173** [289] In "threshold AIS generation" mode, the NE shall insert AIS into the derived DS1 output when the synchronization status message in the
-

OC-N signal that is being used as a reference for that derived DS1 is at or below a user selectable quality level. The default for the threshold shall be quality level 7, "SMC Traceable."

Page 5-83

- R5-174** [290v2] The NE shall validate (as per Section 5.4.7.1) a degradation in the synchronization status message in the OC-N signal to or below the threshold, react as per Section 5.4.5.2.1, and generate AIS (if appropriate) on the derived DS1 timing output within 10 seconds.

Page 5-83

- R5-175** [291] In "message pass-through" mode, the NE shall have the capability to generate the synchronization status messages listed in Table 5-7 in the ESF data link of the derived DS1 signals.

Page 5-83

- R5-176** [292] The derived DS1 output shall carry the synchronization status message that corresponds to the message in the OC-N that is being used as a reference for that derived DS1.

Page 5-83

- R5-177** [293v2] When the synchronization status message in the OC-N used to derive the DS1 changes, the NE shall validate the change (as per Section 5.4.7.1), react as per Section 5.4.5.2.1, and insert the appropriate message in the derived DS1, both within 10 seconds.

Page 5-83

- R5-178** [294] The synchronization status message shall be sent continuously in the ESF data link.

Page 5-83

- CR5-179** [295] An NE may be required to provide a retiming slip buffer for timing distribution on a traffic-carrying, payload DS1 interface.

Page 5-83

- R5-180** [296] The slip buffer shall be at least 1 frame (125 μ s) plus a minimum of 18 μ s of hysteresis. (More hysteresis is desirable.) When a slip occurs, an entire DS1 frame (i.e., 193 bits, including the framing bit) shall be slipped unless the DS1 is byte-synchronously mapped and the VT PTE is generating a new DS1 framing pattern. If a new DS1 framing pattern is being generated, then only the 192 data bits shall be slipped and the framing pattern shall not be affected.

Page 5-84

-
- R5-181** [297] If a retiming slip buffer is provided, the NE shall accumulate slip counts as performance monitoring data according to the criteria in GR-820-CORE, *OTGR Section 5.1: Generic Digital Transmission Surveillance*.
Page 5-84
- R5-182** [298v3] A SONET NE shall meet the criteria related to automatic timing reference switching in Section 3.4 of GR-1244-CORE.
Page 5-84
- R5-183** [300v4] A SONET NE shall support a Manual Reference Switch command that allows the user to manually switch between synchronization references or to the specified reference. This command shall be denied if it would cause a switch to a failed reference or a reference with a lower synchronization status message (including a "DON'T USE for Synchronization" message). It shall also be denied if a Forced Reference Switch command is already active, or if the reference being switched to has been locked out via the Lockout a Reference command (if one or both of those commands are supported). In addition, a Manual Reference Switch shall be preempted if a change occurs or command is received that causes a subsequent timing reference switch (e.g., if the new active reference subsequently fails).
Page 5-85
- CR5-184** [1054] A SONET NE may be required to support a Forced Reference Switch command.
Page 5-86
- R5-185** [1055] If the Forced Reference Switch command is supported, the functionality of that command shall be clearly documented.
Page 5-86
- R5-186** [1056] If the Forced Reference Switch command is supported, a Reference Switch Clear command to clear a Forced Reference Switch shall also be supported.
Page 5-86
- CR5-187** [1057] A SONET NE may be required to support a Lockout a Reference command.
Page 5-86
- R5-188** [1058] If the Lockout a Reference command is supported, its impact shall be to effectively cause the specified reference to be temporarily suspended
-

from the NE's provisioned timing reference list (i.e., until the corresponding Clear command is received).

Page 5-86

- R5-189** [1059] If the Lockout a Reference command is supported, a Reference Lockout Clear command shall also be supported. That command shall cause the NE to clear an existing Lockout a Reference for the specified reference.

Page 5-87

- R5-190** [1060] A SONET NE shall declare (and report to an OS) a standing condition when the first of one or more concurrent or consecutive Forced Reference Switch or Lockout a Reference commands is completed. The level of this condition shall be user-provisionable either as Not Alarmed or as a MN alarm. The condition shall be cleared (and a clear message sent to an OS) when all Forced Reference Switch and/or Lockout a Reference commands have been cleared.

Page 5-87

- R5-191** [1022] An incoming SONET signal shall be considered failed or unavailable for timing purposes under the following conditions:
- ... • Loss of signal energy (e.g., LOS defect detection/failure declaration)
 - ... • Loss of framing (e.g., LOF defect detection/failure declaration)
 - ... • Line AIS (e.g., AIS-L defect detection/failure declaration at an NE containing LTE)

Page 5-87

- R5-192** [302v2] A timing reference shall be considered failed or unavailable upon receipt of synchronization status message indicating that the reference is traceable to a source in holdover/free-run that is of worse quality than the local internal clock.

Page 5-88

- R5-193** [305] If the active timing reference is an OC-N "interface", switching to an alternate timing reference (if available) shall only take place after it has been determined that any available protection switching of the OC-N line and its terminating circuitry has failed to end the outage. The clock shall maintain accuracy and meet the phase transient criteria in Section 5.4.4.3.3 during the protection switch. Traffic and timing need not be protected together.

Page 5-88

-
- R5-194** [310] The available, provisioned reference with the highest quality synchronization status message shall be selected as the active reference.
Page 5-89
- R5-195** [311] The NE shall not select a reference with a "DON'T USE for Synchronization" message as the active synchronization reference.
Page 5-89
- O5-196** [1023] An externally timed NE should refrain for at least X seconds from performing a switch between its primary and secondary external DS1 reference signals based on a change in the synchronization status messages contained in those signals, unless that change causes the NE to consider its currently active reference as failed or unavailable (see **R5-192 [302v2]**), or the message in the active reference changes to the "DON'T USE for Synchronization" message (see **R5-195 [311]**). If a reference switch is still appropriate at the end of this X -second holdoff period, the NE should perform the switch within an additional 1-second period.
Page 5-90
- R5-197** [312] The supplier shall document the scheme used to validate messages on the ESF data link for DS1 synchronization references and on the S1 byte for SONET signals.
Page 5-91
- R5-198** [313] A change in the S1 synchronization status message shall be detected if at least 8 consecutive samples (these may or may not be consecutive SONET frames) of bits 5-8 of the S1 byte have the same (new) value. The sampling rate shall be such that the maximum time to detect a change (assuming no transmission errors) is 1 second.
Page 5-91
- R5-199** [314] A valid synchronization status message in the ESF data link for DS1 synchronization signals shall be detected if at least 7 out of 10 like messages are received.
Page 5-91
- R5-200** [315v2] For S1 messages, if no validated synchronization status message is detected (e.g., due to transmission errors or the receipt of an undefined message) for a period of greater than 10 seconds, the NE shall consider the reference failed.
Page 5-92
-

-
- R5-201** [316] For DS1 references in the SF format, the SONET NE shall consider the reference to have a "Synchronized – Traceability Unknown" message.
Page 5-92
- R5-202** [318] The user shall be able to provision the NE to accept an external DS1 reference in the ESF format that does not support synchronization messages. For such a provisioned reference, the SONET NE shall consider the reference to have a "Synchronized – Traceability Unknown" message.
Page 5-92
- R5-203** [317v2] For DS1 references in the ESF format, if no validated synchronization status message is detected (e.g., due to transmission errors) for a period of greater than 10 seconds, the SONET NE shall consider the reference to be failed (unless the reference has been provisioned as not supporting synchronization status messages).
Page 5-92
- R5-204** [319] The NE shall automatically generate a status report to an OS when the synchronization status message on any provisioned reference changes. The report shall indicate which reference changed, the time of the change, the old synchronization status message, and the new synchronization status message.
Page 5-92
- O5-205** [320v3] The NE should report to the user, on demand, the synchronization status message at any of its SONET "interfaces" (output and input, including those where processing of the incoming S1 byte has been disabled), and on its external DS1 reference signals (when supported).
Page 5-93
- R5-206** [321] When the synchronization status message in the active synchronization reference changes, the NE shall validate the change as per Section 5.4.7.1, react as per Sections 5.4.5.2.1 and 5.4.6.4, and insert the appropriate synchronization status message in all transmitted SONET signals within 10 seconds.
Page 5-93
- R5-207** [322] When the NE enters holdover or free-run, the synchronization status message on all of its transmitted SONET signals shall be changed within 10 seconds to indicate the holdover level of the SONET NE's internal clock (e.g., "Stratum 3 Traceable" or "SMC Traceable").
Page 5-93
-

- R5-208** [323] When the NE recovers from holdover or free-run, the synchronization status message shall not change until the NE has completely resynchronized. The time to change the message shall be no longer than the recovery time [i.e., the sum of the requalification time (as specified in Section 5.4.4.3.5) and the locking time (as specified in GR-1244-CORE)] plus 10 seconds.
Page 5-93
- R5-209** [324] The NE shall allow the user to individually provision each SONET "interface" so that the transmitted signals from that "interface" carry the "DON'T USE for Synchronization" message instead of the message that reflects the actual traceability of the signals.
Page 5-93
- R5-210** [325] If none of the terminating signals at a particular SONET "interface" are being used to derive a DS1, then the NE shall insert the synchronization message from the active external ESF DS1 reference on the SONET signals transmitted from that "interface".
Page 5-94
- R5-211** [326v2] If a terminating signal at a particular SONET "interface" is being used to derive a DS1 and the synchronization status message on the active external ESF DS1 reference matches the synchronization status message on the terminating SONET signal, then the NE shall insert the "DON'T USE for Synchronization" message on the SONET signals transmitted from that "interface" (unless the automatic generation of the DUS message has been disabled for all of the DS1s being derived from that interface, see CR5-214 [1061] and R5-215 [1062]).
Page 5-94
- R5-212** [327v3] If a terminating signal at a particular SONET "interface" is being used to derive a DS1 and (in the steady-state, see R5-213 [1024]) the synchronization status message on the active external ESF DS1 reference does not match the synchronization status message on the terminating SONET signal (or the automatic generation of the DUS message has been disabled for all of the DS1s being derived from that interface, see CR5-214 [1061] and R5-215 [1062]), then the NE shall insert the synchronization message from the active external reference on the SONET signals transmitted from that "interface".
Page 5-94
- R5-213** [1024] When an NE that is transmitting the DUS message at one of its SONET "interfaces" (because it meets the condition described in R5-211 [326v2]) detects a change in the incoming synchronization status message

at that "interface", it shall continue to generate a DUS message for Y seconds after the synchronization status message has been translated to the derived DS1 (i.e., after the message on the outgoing derived DS1 has been changed to reflect the new message received at that SONET "interface").

Page 5-95

CR5-214 [1061] An externally timed SONET NE may be required to be capable of being provisioned such that it does not automatically generate the DUS message according to **R5-211 [326v2]**.

Page 5-95

R5-215 [1062] A SONET NE that can be provisioned to derive each of its DS1s from a different SONET "interface" (see **R5-152 [273v3]**) and that also provides the capability to disable the automatic generation of the DUS message according to **R5-211 [326v2]**, shall provide that capability on a per-derived DS1 basis.

Page 5-95

R5-216 [1063] If a SONET NE provides the capability to disable the automatic generation of the DUS message according to **R5-211 [326v2]**, its default shall be that the automatic generation is enabled.

Page 5-95

R5-217 [328] If the external reference does not carry synchronization messages (e.g., an external DS1 reference is in the SF format or the NE has been provisioned not to expect synchronization messages on the ESF DS1), the NE shall insert the "Synchronized – Traceability Unknown" message on all transmitted SONET signals.

Page 5-96

R5-218 [329] At all SONET "interfaces" that are not the active synchronization reference, the line-timed NE shall insert the synchronization status message from the active synchronization source on the transmitted SONET signals (see Figure 5-24, OC-N #2).

Page 5-96

R5-219 [330] The line-timed NE shall generate a "DON'T USE for Synchronization" message in the transmitted signals from the OC-N "interface" that is the active synchronization reference (see Figure 5-24, OC-N #1).

Page 5-96

-
- R5-220** [331] A through-timed ADM shall insert the synchronization message from the terminating OC-N in the transmitted OC-N in the same direction (see Figure 5-25 for an example).
Page 5-97
- R5-221** [332] A through-timed ADM shall insert the synchronization status message of the appropriate timing source in all dropped SONET signals.
Page 5-97
- R5-222** [333] The framing pattern observed by a SONET NE shall include a subset of the A1 and A2 bytes contained in the incoming STS-N electrical or OC-N signal.
Page 5-98
- R5-223** [334] An SEF defect shall be detected when the incoming signal has a minimum of four consecutive errored framing patterns. The maximum SEF detection time shall be 625 μ s for a random signal.
Page 5-98
- R5-224** [335] The framing algorithm used to check the alignment shall be such that an SEF defect is not detected more than an average of once every 6 minutes while the BER of the STS-N electrical or OC-N signal is 10^{-3} .
Page 5-98
- R5-225** [336] Once an SEF defect has been detected, the SONET NE shall terminate the SEF defect upon detecting two successive error-free framing patterns.
Page 5-98
- R5-226** [337] Timing jitter at network interfaces shall not exceed A_1 Unit Intervals peak-to-peak (UI_{pp}) when measured over a 60-second interval with a bandpass filter having a high-pass cutoff frequency of B_1 (and a roll-off of 20 dB/decade) and a low-pass cutoff frequency of at least B_3 .
... Timing jitter at network interfaces shall not exceed $A_2 UI_{pp}$ when measured over a 60-second interval with a bandpass filter having a high-pass cutoff frequency B_2 (and a roll-off of 20 dB/decade), and a low-pass cutoff frequency of at least B_3 .
Page 5-99
- R5-227** [338] For Category I DS_n interfaces other than DS1 and DS3 (e.g., asynchronous DS1C and DS2 interfaces), the jitter transfer function shall meet the jitter transfer requirement in GR-499-CORE, Section 7.3.2.
Page 5-103
-

-
- R5-228** [339] For Category I DS1 and DS3 interfaces, the jitter transfer function shall be under the mask in Figure 5-26.
Page 5-103
- R5-229** [340] Phase smoothing circuits shall be employed when an asynchronous DS_n payload is demultiplexed from one STS or VT SPE and then multiplexed into another STS or VT SPE within a single SONET NE.
Page 5-103
- R5-230** [341] For Category II interfaces, the jitter transfer function shall be under the mask in Figure 5-27.
Page 5-104
- R5-231** [342] Category I interfaces to SONET NEs shall meet the Category I input jitter tolerance requirement in Section 7.3.1 of GR-499-CORE.
Page 5-105
- R5-232** [343v2] A SONET NE's STS-N electrical and OC-N interfaces, with the exception of OC-48 interfaces that are specified as having "reduced" jitter tolerance (see the discussion below), shall tolerate, as a minimum, input jitter applied according to the mask in Figure 5-28 (with the parameters specified in the figure for the particular rate signal).
Page 5-105
- R5-233** [931] If a SONET NE has reduced OC-48 jitter tolerance, that shall be clearly documented.
Page 5-105
- R5-234** [345] Category II DS1 interfaces to SONET NEs shall meet the Category II input jitter tolerance requirement in Section 7.3.1 of GR-499-CORE.
Page 5-106
- R5-235** [346] For Category I DS_n interfaces other than DS1 and DS3 interfaces, the mapping jitter shall meet the jitter generation requirement in Section 7.3.3 of GR-499-CORE.
Page 5-108
- R5-236** [347] For a Category I DS1 or DS3 interface, the mapping jitter generation shall be less than the value given in Table 5-10.
Page 5-108
- R5-237** [348] Complete data integrity shall be maintained (i.e., no bit errors shall occur) through the SONET system during all of the pointer adjustment
-

jitter generation tests where T is either in the "required range" given in Table 5-11, or is not applicable (i.e., during the single and burst pointer adjustment tests).

Page 5-108

- O5-238** [349] Complete data integrity should be maintained through the SONET system during all of the periodic pointer test sequences where T is in the "objective range" given in Table 5-11.

Page 5-109

- R5-239** [1025] In the presence of a constant frequency offset (between the incoming and outgoing signals at a pointer processor) of up to ± 20 ppm and no pointer adjustments in the incoming signal, the minimum spacing (time) between consecutive pointer adjustments generated by a pointer processor shall be greater than or equal to 0.5 times the long-term average spacing.

Page 5-110

- R5-240** [1026] In the presence of incoming pointer adjustments in the patterns shown in Figure 5-32b (VT only), Figure 5-33b (STS only) or Figure 5-34b (STS and VT) and no frequency offset between the incoming and outgoing signals, the minimum spacing between consecutive pointer adjustments generated by a pointer processor shall be greater than or equal to 0.5 times the long-term average spacing.

Page 5-110

- R5-241** [350] The jitter appearing at a Category I DS_n interface shall be less than the corresponding value in Table 5-12 when the pointer test sequence in Figure 5-29 is applied.

Page 5-111

- R5-242** [351] The jitter appearing at a Category I DS3 interface shall be less than $1.3 UI_{pp}$ when the pointer test sequence described in Figure 5-30 is applied.

Page 5-112

- R5-243** [352] The jitter appearing at a Category I DS3 interface shall be less than $1.2 UI_{pp}$ when the pointer test sequence described in Figure 5-31 is applied.

Page 5-112

- R5-244** [353] The jitter appearing at a Category I DS1 or DS3 interface shall be less than the corresponding value given in Table 5-13 when the pointer test sequences described in Figures 5-32, 5-33, and 5-34 are applied with T in the required range.

Page 5-113

-
- O5-245** [354] The jitter appearing at a Category I DS1 or DS3 interface should be less than the corresponding value given in Table 5-13 when the pointer test sequences described in Figures 5-32, 5-33, and 5-34 are applied with T in the objective range.
Page 5-113
- R5-246** [355] In a single-supplier configuration with a single timing reference offset of 0 to ± 4.6 ppm, the jitter appearing at a Category I DS1 or DS3 interface shall be less than $1.5 U_{I_{pp}}$.
Page 5-114
- O5-247** [356] In a single-supplier configuration with timing reference offsets equal to twice the specified free-run accuracy of the NEs' internal clocks (e.g., 9.2 ppm for NEs with stratum 3 clocks, 40 ppm for NEs with SMCs), the jitter appearing at a Category I DS1 or DS3 interface should be less than $1.5 U_{I_{pp}}$.
Page 5-114
- R5-248** [357] The jitter generated at Category II interfaces shall be less than $0.01 U_{I_{rms}}$, and shall also be less than $0.10 U_{I_{pp}}$.
Page 5-114
- R5-249** [932] The jitter appearing at the output of a "nominal" DS1 or DS3 desynchronizer shall be less than the corresponding value in Table 5-12 when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in Figure 5-29.
Page 5-115
- R5-250** [933] The jitter appearing at the output of a "nominal" DS3 desynchronizer shall be less than $1.3 U_{I_{pp}}$ when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in Figure 5-30.
Page 5-115
- R5-251** [934] The jitter appearing at the output of a "nominal" DS3 desynchronizer shall be less than $1.2 U_{I_{pp}}$ when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in Figure 5-32.
Page 5-115
- R5-252** [935] The jitter appearing at the output of a "nominal" DS1 or DS3 desynchronizer shall be less than the corresponding value given in Table 5-13 when that desynchronizer receives a signal from the NE, and
-

the input signal to the NE contains pointer adjustments in the pointer test sequences shown in Figures 5-32, 5-33, and 5-34, with T in the required range.

Page 5-115

- O5-253** [936] The jitter appearing at the output of a "nominal" DS1 or DS3 desynchronizer should be less than the corresponding value given in Table 5-13 when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequences shown in Figures 5-32, 5-33, and 5-34, with T in the objective range.

Page 5-115

- R5-254** [358] The mapping phase variations on a DS1 output from a SONET NE shall be below the mask in Figure 5-35.

Page 5-119

- R5-255** [1027] The mapping phase variations on a DS3 output from a SONET NE shall be below the mask in Figure 5-36.

Page 5-120

- R5-256** [359] The MTIE of a DS1 output from a SONET NE shall be below the mask in Figure 5-37 when the pointer adjustment test sequence in Figure 5-29 is applied.

Page 5-122

- R5-257** [1028] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-38 when the pointer adjustment test sequence in Figure 5-29 is applied.

Page 5-122

- R5-258** [1029] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-39 when the pointer test sequence described in Figure 5-30 is applied.

Page 5-124

- R5-259** [1030] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-40 when the pointer test sequence described in Figure 5-31 is applied.

Page 5-125

- R5-260** [360v3] The MTIE of a DS1 output from a SONET NE shall be below the mask of Figure 5-41 when the pointer adjustment test sequences in

Figures 5-32b and 5-34b are applied with T in the required range (see Table 5-11).

Page 5-126

O5-261 [937v2] The MTIE of a DS1 output from a SONET NE should be below the mask of Figure 5-41 when the pointer adjustment test sequences in Figures 5-32b and 5-34b are applied with T in the objective range.

Page 5-127

R5-262 [1031v2] The MTIE of a DS3 output from a SONET NE shall be below the mask of Figure 5-42 when the pointer adjustment test sequences in Figures 5-33b and 5-34b are applied with T in the required range.

Page 5-127

O5-263 [1032v2] The MTIE of a DS3 output from a SONET NE should be below the mask of Figure 5-42 when the pointer adjustment test sequences in Figures 5-33b and 5-34b are applied with T in the objective range.

Page 5-127

R6-1 [366] Any provisionable feature or parameter shall be provisionable locally by a craftsperson and remotely from an OS.

Page 6-2

R6-2 [361v2] A SONET NE shall use the two-level "Group #, VT #" convention shown in Section 3 of this document (Figures 3-11, 3-13, 3-15, and 3-17) for numbering VTs within an STS-1. This numbering convention shall be used on OS/NE, WS/NE (including any GUI display), and NE/NE interfaces.

Page 6-2

R6-3 [938] A SONET NE shall use either the two-level "STS-3 #, STS-1 #" convention or the single-level "1 to N in order of appearance at the input to the byte-interleaver" convention for numbering STS-1s within an OC-N. These numbering conventions shall be used on OS/NE, WS/NE (including any GUI display), and NE/NE interfaces.

Page 6-2

R6-4 [939] For numbering an STS-Mc within an OC-N, a SONET NE shall use the number of the STS-1 in which the STS-Mc starts. This numbering convention shall be used on OS/NE, WS/NE (including any GUI display), and NE/NE interfaces.

Page 6-4

-
- R6-5** [367] A SONET NE shall provide, via the local craftsperson interface, the ability to initialize the NE with communications-related information upon installation of the NE by the network provider. Such information shall include protocol options for the various operations communications interfaces supported by the NE, the NE's TID (when TL1 messages are supported), and the NE's network address (NSAP) for communications purposes. The supplier shall clearly specify in user documentation those data items that need to be provided upon installation of the NE to ensure proper operations communications involving the NE.
Page 6-4
- R6-6** [368] The LTE shall be provisionable to indicate either physical-layer regenerators or section-terminating regenerators are being used for a given line.
Page 6-4
- R6-7** [369] A SONET NE shall provide a local, primary, nonvolatile memory backup.
Page 6-5
- R6-8** [370] Data shall be backed up in at least one nonvolatile backup memory automatically after each primary memory update.
Page 6-5
- R6-9** [371] Restoration of data from the local backup memory, once initiated, shall be completed within 5 minutes.
Page 6-5
- O6-10** [372] Restoration of data from the local backup memory, once initiated, should take no more than one minute.
Page 6-5
- R6-11** [373] A SONET NE shall be able to have its configurable memory restored by a remote memory restoration application identified by the network provider.
Page 6-5
- R6-12** [374] A SONET NE shall be able to determine whether or not the source of an update to its configurable memory is the same management application (e.g., OS) as the one responsible for restoring lost primary (and secondary) nonvolatile memory backups.
Page 6-5
-

-
- R6-13** [375] When the source of a memory update is different from the memory restoration application (e.g., a local craftsperson or other management application), the SONET NE shall send an autonomous indication to the memory restoration application detailing this "hidden update."
Page 6-5
- R6-14** [376] If communications to the memory restoration application are not available, "hidden updates" shall be logged by the SONET NE and reported when asked for by the memory restoration application after communications are restored.
Page 6-5
- CR6-15** [377] A SONET NE may be required to support the ability to have its nonvolatile memory backup restored via bulk file transfer methods by a remote memory restoration management application (e.g., an OS or controller). This feature may be considered an objective or a requirement for certain types of SONET NEs, and suppliers are referred to NE-specific requirements for such instances.
Page 6-6
- R6-16** [378] If a SONET NE supports the optional bulk memory restoration feature, then the NE shall support the feature via a full 7-layer OSI-based operations interface to a bulk memory restoration application using the memory backup functions and FTAM protocol requirements described in GR-1250-CORE, *Generic Requirements for Synchronous Optical Network (SONET) File Transfer*.
Page 6-6
- R6-17** [379] If a SONET NE supports the optional bulk memory restoration feature, then the NE shall be able to report a bulk "snapshot" of its nonvolatile memory backup to the memory restoration application upon request.
Page 6-6
- R6-18** [380] A security mechanism shall be provided within a SONET NE to prevent unauthorized communication to the NE via any ports and communications channels accepting operations-related command inputs, and to allow secure access to the database of the NE. Such a mechanism shall adhere to the security requirements of GR-815-CORE.
Page 6-7
-

-
- R6-19** [381] The data necessary to support the security mechanism within a SONET NE shall be provided and administered only by authorized security administrators via either WS/NE or OS/NE communications.
Page 6-7
- R6-20** [382] A SONET NE shall support security administration functions in conformance with GR-815-CORE.
Page 6-7
- R6-21** [383] A SONET NE supporting interfaces conforming to 7-layer OSI protocol stacks shall conform to the security requirements described in GR-1469-CORE, for layers 1 through 6.
- ... • When TL1-based interfaces are used in the NE for WS/NE, NE/NE, or OS/NE communications, the NE shall support the data and messages provided in TR-NWT-000835 for the administration of the NE security mechanism.
 - ... • When CMISE interfaces are used on the application layer, the NE shall support data, messages, and mechanisms to conform to the requirements provided by GR-1469-CORE.
- Page 6-8
- R6-22** [384] The data communications network component of the SONET TMN shall conform to the security requirements in GR-1332-CORE.
Page 6-9
- R6-23** [385] A SONET NE shall support identification and authentication for all users, for all ports accepting operations-related command inputs, in conformance with GR-815-CORE.
Page 6-9
- R6-24** [386] A SONET NE which supports remote access for operations-related command inputs shall provide a feature for additional strong authentication, beyond reusable passwords, such as:
- ... • Third-party authentication
 - ... • Public/private key encryption technology
- Page 6-9
- R6-25** [387] When TL1 interfaces are used, a SONET NE shall enforce that a session requester accessing the NE via a TL1/X.25-based OS/NE interface
-

must pass identification information based on X.25 calling address or PVC identifier.

Page 6-9

- R6-26** [388] A SONET NE shall support system access control functions, in conformance with GR-815-CORE.

Page 6-9

- R6-27** [389] A SONET NE shall not grant a user remote access unless the user is authenticated via strong authentication.

Page 6-9

- R6-28** [390] A SONET NE shall employ features corresponding to the timeout interval function that GR-815-CORE describes.

Page 6-10

- R6-29** [391] A SONET NE shall break down an OSI Application Association if an attempted session request is unsuccessful after a provisionable number of tries. The default number of tries shall be three.

Page 6-10

- R6-30** [392] A SONET NE shall support resource access control and authorization functions in conformance with GR-815-CORE.

Page 6-10

- R6-31** [393v2] A SONET NE supporting one or more restricted DCCs shall support user identification and access control privileges (see GR-815-CORE) that limit the functionality available to valid outside users accessing the NE via a restricted DCC. The functions allowed via such DCCs shall be definable by local service providers.

Page 6-10

- R6-32** [394] A SONET NE shall support audit features, in conformance with GR-815-CORE.

Page 6-10

- R6-33** [395] An NE shall be capable of disabling a Section DCC. The default shall be to enable an equipped DCC.

Page 6-11

- R6-34** [397] Only properly authorized system administrators shall be allowed to enable or disable an equipped Section DCC.

Page 6-11

-
- R6-35** [398] On reinitialization of the NE after failure, DCCs shall maintain the enabled or disabled state they held before the failure of the NE.
Page 6-11
- CR6-36** [399] To provide NE/NE and indirect OS/NE communications paths that may be required by a network provider across administrative boundaries, an NE may be required to terminate one or more enabled DCCs that cross an administrative boundary.
Page 6-11
- R6-37** [400] If an NE has the capability to terminate a DCC that crosses an administrative boundary, it shall be capable of classifying it as either restricted or unrestricted for operations security purposes. The default setting shall be unrestricted.
Page 6-11
- R6-38** [402] Only properly authorized system administrators shall be allowed to change the classification of a DCC from restricted to unrestricted or vice versa.
Page 6-11
- R6-39** [403] A change in classification of a restricted DCC shall not be allowed by communications over that or any other restricted DCC.
Page 6-11
- R6-40** [404] On reinitialization of the NE after failure, DCCs shall maintain the restricted or unrestricted states they held before the failure of the NE.
Page 6-11
- R6-41** [405] A SONET NE shall not forward received NPDUs across an administrative boundary to the underlying Data Link Service for a restricted DCC. The NE shall "terminate" the DCC.
Page 6-12
- R6-42** [406] A SONET NE shall support a list for each restricted DCC, that itemizes NSAP address pairs that are to be allowed in source and destination address fields of NPDUs allowed into the NE's network.
Page 6-12
- R6-43** [407] The list described in **R6-42 [406]** shall be established via either WS/NE or OS/NE communications, and only by authorized system administrators via an unrestricted DCC or direct interfaces.
Page 6-12
-

- R6-44** [408] A SONET NE shall screen the NPDUs received from a restricted DCC and only accept the PDU if the source address and destination address matches an allowable source/destination address pair from the list described in **R6-42 [406]**.
Page 6-12
- R6-45** [409] A SONET NE shall enforce access control on the restricted DCC to restrict the requested operations functions that will be permitted on a per-user basis.
Page 6-12
- R6-46** [410] A SONET NE shall provide the ability for an authorized administrator to specify the access control privileges assigned to a user for restricted DCC use.
Page 6-12
- R6-47** [411] The initial software generic shall be entered in the SONET NE at or before installation.
Page 6-12
- O6-48** [412] A SONET NE should be able to receive its initial software load and later software generics via either an OS/NE or NE/NE (or MD/NE) interface using FTAM.
Page 6-12
- R6-49** [413] The NE shall provide the ability to retrieve (locally via a WS and remotely via an OS) the current version ID of software.
Page 6-13
- O6-50** [414] Software updates, or patches, should be identified and included in the current version report.
Page 6-13
- O6-51** [415] A SONET NE should be able to report to a management application or a craftsperson its equipage (including plug-ins, common equipment, and software), option settings, and its crossconnect configuration.
Page 6-13
- R6-52** [416v2] A SONET NE shall monitor all incoming SONET signals (before descrambling) for an "all-zeros patterns," where an all-zeros pattern corresponds to no light pulses for OC-N optical interfaces and no voltage transitions for STS-1 and STS-3 electrical interfaces. An LOS defect shall be detected when an all-zeros pattern on the incoming SONET signal
-

lasts 100 μ s or longer. If an all-zeros pattern lasts 2.3 μ s or less, an LOS defect shall not be detected.

Page 6-16

- O6-53** [940] If a SONET NE monitors the received signal level for the purpose of detecting LOS defects, then its signal level threshold should be selected such that an LOS defect will not be detected if the BER is still acceptable (e.g., no LOS defect if the BER is better than the SF BER threshold used for protection switching in linear APS, see Section 5.3.3.1).

Page 6-16

- R6-54** [417] The SONET NE shall terminate an LOS defect when the incoming signal has two consecutive valid framing alignment patterns and, during the intervening time (one frame), no all-zeros pattern qualifying as an LOS defect exists.

Page 6-16

- R6-55** [418v2] A SONET NE shall declare an LOS failure when the LOS defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an LOS failure indication and send an alarm message to an OS.

Page 6-16

- R6-56** [419] For the purposes of trunk conditioning, a SONET NE that contains DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall declare an LOS failure if it is subject to a period of short, intermittent LOS defects. For failures resulting from the NE intermittently detecting and terminating the defect, the NE shall integrate the time during which the defect persists, using a 4:1 to 15:1 count-up/count-down ratio. During a sustained LOS defect, the integrator shall count up to reach the alarm threshold in 2.5 (± 0.5) seconds. Upon reaching the alarm threshold, the NE shall declare the LOS failure, set an LOS failure indication, and send an alarm message to an OS. If the defect is terminated before the threshold is reached, the integrator shall count down at a slope 1/4 to 1/15 of the count-up slope.

Page 6-16

- R6-57** [420] A SONET NE shall clear an LOS failure when the LOS defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the SONET NE shall clear the LOS failure indication and send a clear message to an OS.

Page 6-17

- R6-58** [421v2] SONET NEs interfacing with DS1, DS1C, DS2, or DS3 signals shall detect LOS on those signals according to the requirements in GR-499-CORE.
Page 6-17
- R6-59** [422] All incoming SONET signals shall be monitored for LOF. A SONET NE shall detect an LOF defect when an SEF defect (see Section 5.5) on the incoming SONET signal persists for 3 ms.
Page 6-17
- R6-60** [423] If the optional integration timer described above is provided for LOF monitoring, the supplier shall clearly describe its use in the product documentation.
Page 6-18
- R6-61** [424] The SONET NE shall terminate an LOF defect 1 ms to 3 ms after terminating the SEF defect on the incoming SONET signal, if the SEF defect is not (re)detected before the LOF defect is terminated.
Page 6-18
- O6-62** [425] The SONET NE should terminate an LOF defect 1 ms after terminating the SEF defect on the incoming SONET signal, if the SEF defect is not detected within the 1-ms time period. (This objective is not applicable if the optional 3-ms integration timer, described above, is used.)
Page 6-18
- R6-63** [426v2] A SONET NE shall declare an LOF failure when an LOF defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOF failure, the NE shall set an LOF failure indication and send an alarm message to an OS unless the condition in **R6-285** [626v2] applies (see Section 6.2.1.8).
Page 6-18
- R6-64** [427] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56** [419] to declare LOF failures. Upon declaring an LOF failure, the NE shall perform the actions listed in **R6-63** [426v2].
Page 6-18
- R6-65** [428v2] A SONET NE shall clear an LOF failure when the LOF defect is absent for 10 (± 0.5) seconds. Upon clearing the LOF failure, the SONET

NE shall clear the LOF failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-18

R6-66 [429] An NE shall monitor for DS_n OOF on DS_n paths that are terminated by the NE.

Page 6-18

R6-67 [941] If an NE supports an asynchronous DS_n mapping as defined in Section 3.4, then it shall be capable of providing clear-channel transport of DS_n signals using that mapping.

Page 6-19

CR6-68 [942] An NE that supports an asynchronous DS_n mapping may be required to provide non-clear-channel transport where the incoming DS_n signal is monitored for DS_n OOF.

Page 6-19

CR6-69 [943] An NE that supports an asynchronous DS_n mapping may be required to provide non-clear-channel transport where the outgoing DS_n signal (i.e., the DS_n signal that is demultiplexed from the STS or VT SPE) is monitored for DS_n OOF.

Page 6-19

R6-70 [430v3] STS PTE and LTE that processes the STS pointers shall monitor for LOP-P. An LOP-P defect shall be detected if a valid pointer is not found in N consecutive frames (where $8 \leq N \leq 10$), or if N consecutive NDFs (other than in a concatenation indicator, see Section 3.5.1.3) are detected. An LOP-P defect shall not be detected when LTE is receiving and relaying an all-ones STS pointer, or when STS PTE is receiving pointers that qualify as those necessary to cause the detection of an AIS-P defect (i.e., three or more consecutive all-ones pointers).

Page 6-20

R6-71 [944v2] VT PTE and STS PTE that processes VT pointers shall monitor for LOP-V. An LOP-V defect shall be detected if a valid pointer is not found in N consecutive superframes (where $8 \leq N \leq 10$), or if N consecutive NDFs are detected. An LOP-V defect shall not be detected when STS PTE is receiving and relaying an all-ones VT pointer, or when VT PTE is receiving pointers that qualify as those necessary to cause the detection of an AIS-V defect.

Page 6-20

-
- R6-72** [431v2] STS PTE and LTE that processes the STS pointers shall terminate an LOP-P defect when the STS has a valid pointer with a normal NDF, or a valid concatenation indicator, in three consecutive frames.
Page 6-22
- R6-73** [432v2] STS PTE shall terminate an LOP-P defect when it detects an AIS-P defect.
Page 6-22
- R6-74** [945] LTE that processes STS pointers shall terminate an LOP-P defect when it relays an all-ones STS pointer.
Page 6-22
- R6-75** [946] VT PTE and STS PTE that processes VT pointers shall terminate an LOP-V defect when the VT has a valid pointer with a normal NDF in three consecutive superframes.
Page 6-22
- R6-76** [947] VT PTE shall terminate an LOP-V defect when it detects an AIS-V defect.
Page 6-22
- R6-77** [948] STS PTE that processes VT pointers shall terminate an LOP-V defect when it relays an all-ones VT pointer.
Page 6-22
- R6-78** [433v2] A SONET NE shall declare an LOP-P failure when an LOP-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOP-P failure, the NE shall set an LOP-P failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
Page 6-22
- R6-79** [949] A SONET NE shall declare an LOP-V failure when an LOP-V defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOP-V failure, the NE shall set an LOP-V failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
Page 6-22
- R6-80** [434] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare LOP-P and LOP-V failures. Upon
-

declaring an LOP-P or LOP-V failure, the NE shall perform the actions listed in R6-78 [433v2] or R6-79 [949].

Page 6-22

R6-81 [435v2] A SONET NE shall clear an LOP-P failure when an LOP-P defect is absent for 10.0 (± 0.5) seconds. Upon clearing the LOP-P failure, the SONET NE shall clear the LOP-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-22

R6-82 [950] A SONET NE shall clear an LOP-V failure when an LOP-V defect is absent for 10.0 (± 0.5) seconds. Upon clearing the LOP-V failure, the SONET NE shall clear the LOP-V failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-22

R6-83 [436] Equipment failures shall be classified as either Service-Affecting (SA) or Non-Service-Affecting (NSA), depending on whether they affect the services that the equipment transports.

Page 6-23

R6-84 [437] Equipment failures shall be classified as critical, major, or minor.

Page 6-23

R6-85 [438] Because hardware designs vary, the report of the equipment failure shall describe the failure condition.

Page 6-23

R6-86 [439] The NE shall be able to declare the following equipment failures (as a minimum):

- ... • Fuse or power circuit failures
- ... • Synchronization equipment failures
- ... • Protection switching equipment failures
- ... • CPU failures
- ... • Local nonvolatile backup memory failures
- ... • SONET signal origination and termination equipment failures
- ... • Receiver failures (optical detector failures)
- ... • Transmitter failures (light source failure, including laser failures)

- ... • Non-SONET signal (e.g., DS_n) origination and termination equipment failures
- ... • Switching matrix module failures (if cross-connect functionality is provided)
- ... • DCC hardware failures (also see Section 6.2.1.1.7)
- ... • Manual removal of in-service (i.e., active) equipment.

Page 6-23

R6-87 [440] Upon declaring an equipment failure, a SONET NE shall

- ... • Switch to duplex or standby equipment, if available
- ... • Set a local indication
- ... • Send an alarm message to an OS.

Page 6-23

R6-88 [441] Upon clearing an equipment failure, a SONET NE shall clear the equipment failure indication and send a clear message to the OS.

Page 6-24

CR6-89 [442] A SONET NE may be required to detect and report certain environmental conditions in some applications (e.g., NEs in a CEV).

Page 6-24

O6-90 [443] A SONET NE should detect operating system or other software errors, and report them to an OS, independently of CPU hardware failures.

Page 6-24

R6-91 [444] A SONET NE shall have the capability of declaring a Loss of Synchronization failure, due to either loss of primary timing reference or loss of secondary timing reference. Upon declaring the failure, the NE shall set a Loss of Synchronization failure indication and send a message to an OS. The message shall include an indication of reference switches and the reason for failure (e.g., LOS, LOF or OOF, synchronization message, etc.).

Page 6-24

R6-92 [445] Upon clearing the Loss of Synchronization failure, a SONET NE shall clear the Loss of Synchronization failure indication and send a clear message to the appropriate OS.

Page 6-24

-
- R6-93** [446] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Protection Switching Byte defect within 50 ms of the occurrence of either an inconsistent APS byte or an invalid code, unless the condition for terminating the defect occurs.
Page 6-25
- R6-94** [447] LTE shall not detect a Protection Switching Byte defect when it has detected an AIS-L defect.
Page 6-26
- R6-95** [448] LTE shall terminate the Protection Switching Byte defect within 50 ms of the occurrence of three consecutive, identical, and valid APS codes, unless the condition for detecting the defect occurs.
Page 6-26
- R6-96** [449] LTE shall not terminate a Protection Switching Byte defect when it has detected the AIS-L defect.
Page 6-26
- R6-97** [450] An NE shall declare a Protection Switching Byte failure when a Protection Switching Byte defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall perform the following actions:
- ... 1. A Protection Switching Byte failure indication shall be set and an alarm message shall be sent to an OS.
 - ... 2. If a working channel that was being selected from the protection line is switched back to the working line (see Section 5.3.5.5), a message shall be sent to an OS indicating the switch back to the working line.
- Page 6-26
- R6-98** [451] An NE shall clear the Protection Switching Byte failure when the Protection Switching Byte defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Protection Switching Byte failure indication and send an alarm clear message to an OS.
Page 6-26
- R6-99** [452] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte do not match for 50 ms.
Page 6-26
-

-
- O6-100** [453] LTE operating in a linear APS mode other than the 1+1 unidirectional mode should detect a Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte are mismatched in three consecutive frames.
Page 6-27
- R6-101** [454] LTE shall not detect a Channel Mismatch defect when it has detected an AIS-L defect.
Page 6-27
- R6-102** [455] LTE shall terminate the Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte match in three consecutive frames.
Page 6-27
- R6-103** [456] LTE shall not terminate a Channel Mismatch defect when it has detected an AIS-L defect.
Page 6-27
- R6-104** [457] An NE shall declare a Channel Mismatch failure if the Channel Mismatch defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a Channel Mismatch failure indication and send an alarm message to an OS.
Page 6-27
- R6-105** [458] An NE shall clear the Channel Mismatch failure if the Channel Mismatch defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Channel Mismatch failure indication and send an alarm clear message to the OS.
Page 6-27
- R6-106** [459] LTE provisioned to operate in a linear APS mode other than the 1+1 unidirectional mode shall detect an APS Mode Mismatch defect within 100 ms of receiving the first of five consecutive samples of frames (which may or may not be consecutive frames) with identical mode information (either in bit 5 of K2 or bits 6-8 of K2) that is mismatched, as defined above, unless the condition for terminating the defect occurs before the defect is detected.
Page 6-28
- R6-107** [460] LTE shall not detect an APS Mode Mismatch defect when it has detected an AIS-L defect.
Page 6-28
-

-
- R6-108** [461] LTE shall terminate an APS Mode Mismatch defect within 50 ms of receiving the first of five consecutive samples of frames (which may or may not be consecutive frames) with identical mode information that is not mismatched as defined above, unless the condition for detecting the defect occurs before terminating the defect.
Page 6-28
- R6-109** [462] LTE shall not terminate an APS Mode Mismatch defect when it has detected an AIS-L defect.
Page 6-28
- R6-110** [463] An NE shall declare an APS Mode Mismatch failure when an APS Mode Mismatch defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set the APS Mode Mismatch failure indication and send an alarm message to an OS.
Page 6-28
- R6-111** [464] An NE shall clear the APS Mode Mismatch failure if the APS Mode Mismatch defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the APS Mode Mismatch failure indication and send an alarm clear message to the OS.
Page 6-28
- R6-112** [465] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Far-End Protection-Line defect when it receives three consecutive K1 bytes with the code indicating "SF on the protection line."
Page 6-29
- R6-113** [466] LTE shall terminate the Far-End Protection-Line defect when it receives three consecutive, identical, and valid K1 bytes with any code other than "SF on the protection line."
Page 6-29
- R6-114** [467v2] An NE shall declare a Far-End Protection-Line failure when a Far-end Protection-Line defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, the NE shall set a Far-End Protection-Line failure indication and (if it is a Reported failure) send an alarm message to an OS.
Page 6-29
- R6-115** [468] An NE shall clear the Far-End Protection-Line failure if the Far-End Protection-Line defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Far-End Protection-Line failure indication and
-

send an alarm clear message to the OS (if the failure was reported to the OS).

Page 6-29

R6-116 [1064] When a BER-based SF condition persists for 2.5 (± 0.5) seconds, an NE shall set an SF BER indication and send an alarm message to an OS.
Page 6-29

R6-117 [1065] When a BER-based SD condition persists for 2.5 (± 0.5) seconds, an NE shall set an SD BER indication and send an alarm message to an OS.
Page 6-29

R6-118 [1066] An NE shall clear an SF BER indication 10 (± 0.5) seconds after detecting that the BER is less than the SF clearing threshold, assuming that it does not detect that the BER is greater than the SF threshold during that time period (see Section 5.3.4). Upon clearing an SF BER indication, the NE shall send an alarm clear message to the OS (if the SF BER was reported to the OS).

Page 6-29

R6-119 [1067] An NE shall clear an SD BER indication 10 (± 0.5) seconds after detecting that the BER is less than the SD clearing threshold, assuming that it does not detect that the BER is greater than the SD threshold during that time period (see Section 5.3.4). Upon clearing an SD BER indication, the NE shall send an alarm clear message to the OS (if the SD BER was reported to the OS).

Page 6-30

R6-120 [1068] As a default, the alarm level for SF BER and SD BER indications shall be Not Alarmed.

Page 6-30

R6-121 [469] If LTE that provides the linear APS feature is unable to perform automatic protection switching upon receiving an AIS-L signal, it shall send an SA alarm to an OS, indicating the inability to perform a protection switch.

Page 6-30

R6-122 [470] A SONET NE shall be able to declare a DCC failure. Upon declaring the failure, it shall set a DCC failure indication and send a message to an OS.

Page 6-30

-
- R6-123** [471] Upon clearing a DCC failure, a SONET NE shall clear the DCC failure indication and send a clear message to an OS. Page 6-30
- R6-124** [472] STS PTE shall detect an STS Payload Label Mismatch (PLM-P) defect within 250 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of mismatched STS Signal Labels (C2 byte), as specified in Table 6-2. Page 6-31
- O6-125** [473] STS PTE should detect a PLM-P defect immediately upon receipt of five contiguous frames with mismatched STS Signal Labels, as specified in Table 6-2. Page 6-31
- R6-126** [476] STS PTE shall terminate a PLM-P defect within 250 ms of detecting the onset of at least five consecutive samples (which may or may not be consecutive frames) of matched STS Signal Labels, as specified in Table 6-2. Page 6-31
- O6-127** [477] STS PTE should terminate a PLM-P defect immediately upon receipt of five contiguous frames with matched STS Signal Labels, as specified in Table 6-2. Page 6-32
- R6-128** [478v2] STS PTE shall terminate a PLM-P defect upon detecting an UNEQ-P defect. Page 6-32
- R6-129** [479] An NE shall declare a PLM-P failure if a PLM-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a PLM-P failure indication and send an alarm message to an OS unless the condition in R6-285 [626v2] applies. Page 6-32
- R6-130** [480v2] An NE shall clear the PLM-P failure if the PLM-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the PLM-P failure indication and send a clear message to an OS (if the failure was reported to the OS). Page 6-32
- R6-131** [481] STS PTE shall detect an STS Path Unequipped (UNEQ-P) defect within 10 ms of the onset of at least five consecutive samples (which may
-

or may not be consecutive frames) of unequipped STS Signal Labels (C2 byte), as specified in Table 6-2.

Page 6-32

- O6-132** [482] STS PTE should detect an UNEQ-P defect immediately upon receipt of five contiguous frames with unequipped STS Signal Labels, as specified in Table 6-2.

Page 6-32

- R6-133** [485] STS PTE shall terminate an UNEQ-P defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of STS Signal Labels that are not unequipped, as specified in Table 6-2.

Page 6-32

- O6-134** [486] STS PTE should terminate an UNEQ-P defect immediately upon receipt of five contiguous frames with STS Signal Labels that are not unequipped, as specified in Table 6-2.

Page 6-32

- R6-135** [488] An NE shall declare an UNEQ-P failure if an UNEQ-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an UNEQ-P failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.

Page 6-32

- R6-136** [489v2] An NE shall clear the UNEQ-P failure if the UNEQ-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the UNEQ-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-33

- R6-137** [491] VT PTE shall detect a VT Payload Label Mismatch (PLM-V) defect within 250 ms of the onset of at least five consecutive samples (which may or may not be consecutive superframes) of mismatched VT Signal Labels, as specified in Table 6-3.

Page 6-34

- O6-138** [492] VT PTE should detect a PLM-V defect immediately upon receipt of five contiguous superframes with mismatched VT Signal Labels, as specified in Table 6-3.

Page 6-35

-
- R6-139** [495] VT PTE shall terminate a PLM-V defect within 250 ms of detecting the onset of at least five consecutive samples (which may or may not be consecutive superframes) of matched VT Signal Labels, as specified in Table 6-3.
Page 6-35
- O6-140** [496] VT PTE should terminate a PLM-V defect immediately upon receipt of five contiguous superframes with matched VT Signal Labels, as specified in Table 6-3.
Page 6-35
- R6-141** [497v2] VT PTE shall terminate the PLM-V defect immediately upon detecting an UNEQ-V defect on the incoming signal.
Page 6-35
- R6-142** [498] An NE shall declare a PLM-V failure when a PLM-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a PLM-V failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
Page 6-35
- R6-143** [499] An NE shall clear a PLM-V failure when the PLM-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the PLM-V failure indication and send a clear message to the OS (if the failure was reported to the OS).
Page 6-35
- R6-144** [501] VT PTE shall detect a VT Path Unequipped (UNEQ-V) defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive superframes) of unequipped VT Signal Labels, as specified in Table 6-3.
Page 6-35
- O6-145** [502] VT PTE should detect an UNEQ-V defect immediately upon receipt of five contiguous superframes with unequipped VT Signal Labels, as specified in Table 6-3.
Page 6-36
- R6-146** [505] VT PTE shall terminate an UNEQ-V defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive superframes) of VT Signal Labels that are not unequipped, as specified in Table 6-3.
Page 6-36
-

- O6-147** [506] VT PTE should terminate an UNEQ-V defect immediately upon receipt of five contiguous superframes with VT Signal Labels that are not unequipped, as specified in Table 6-3.
Page 6-36
- R6-148** [508] An NE shall declare an UNEQ-V failure when an UNEQ-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an UNEQ-V failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.
Page 6-36
- R6-149** [509] An NE shall clear an UNEQ-V failure when the UNEQ-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the UNEQ-V failure indication and send a clear message to the OS (if the failure was reported to the OS).
Page 6-36
- R6-150** [725v2] A SONET NE that contains STS PTE shall allow the user to provision, on a per-path basis, the contents of the STS Path Trace carried in the J1 byte of the STS path overhead originated by the PTE. The transmitted STS Path Trace string shall be 64 bytes in length.
Page 6-38
- R6-151** [997v2] A SONET NE that contains STS PTE shall support a feature that allows the contents of the STS Path Traces to be provisioned as ASCII characters. In addition, the following apply:
- ... • The feature shall allow the user to enter a string of up to 62 characters
 - ... • The feature shall place no restriction on the content of the string, except that the characters shall be ASCII printable characters
 - ... • The NE shall automatically pad the string entered by the user to 62 characters using ASCII NULL characters, and then add <CR> and <LF> characters (i.e., '0D' and '0A') for a total of 64 characters.
 - ... • Each 8-bit ASCII character shall be loaded into one J1 byte.
- Page 6-38
- R6-152** [1069] A SONET NE shall support a feature to allow the user to provision the "expected" ASCII-based path trace for each STS path that it terminates and for which TIM-P detection has been activated (see below). In addition, the following apply:
- ... • The feature shall allow the user to enter a string of up to 62 characters

- The feature shall place no restriction on the contents of the string, except that the characters shall be ASCII printable characters.

Page 6-39

- R6-153** [1070] A SONET NE that contains STS PTE shall support a feature that, when activated monitors the received STS path trace string for TIM-P detection purposes. That feature shall be provisionable on a per-STs path basis, and shall have a default of "not active".

Page 6-39

- R6-154** [1071] STS PTE shall detect a TIM-P defect within 30 seconds (or less) when none of the sampled 64-byte STS path trace strings match the provisioned expected value.

Page 6-40

- R6-155** [1072] A SONET NE's TIM-P defect detection algorithm shall be such that, given an incoming signal with a BER of 10^{-3} or less, the probability that the STS PTE will detect a false TIM-P defect during the TIM-P defect detection time is less than 1×10^{-n} .

Page 6-40

- R6-156** [1073] A change in the phase of an incoming STS path trace string shall cause the STS PTE to consider, at most, one sample to be mismatched for the purpose of detecting and terminating TIM-P defects.

Page 6-41

- O6-157** [1001v2] If a SONET NE is comparing an incoming path trace string to an expected path trace (for either TIM-P defect detection/termination or diagnostics purposes) and the expected path trace consists of ASCII characters, then the comparison should ignore any trailing ASCII NULL, <CR>, and <LF> characters contained in the incoming path trace.

Page 6-41

- O6-158** [1074] A SONET NE's STS path trace sampling rate and TIM-P defect detection algorithm should be such that a TIM-P defect is detected within 2 seconds after the NE's STS PTE stops receiving the provisioned expected path trace string.

Page 6-41

- R6-159** [1075] STS PTE shall terminate a TIM-P defect within 30 seconds (or less) when four-fifths (or more) of the sampled STS path trace strings match the provisioned expected value.

Page 6-41

O6-160 [1076] A SONET NE's STS path trace sampling rate and TIM-P defect termination algorithm should be such that a TIM-P defect is terminated within 2 seconds after the STS PTE begins receiving the provisioned expected path trace string.

Page 6-41

R6-161 [1077] An NE shall declare a TIM-P failure if a TIM-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a TIM-P failure indication and send an alarm message to an OS unless the condition in **R6-285 [626v2]** applies.

Page 6-41

R6-162 [1078] An NE shall clear the TIM-P failure if the TIM-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the TIM-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-42

R6-163 [512v2] STE shall generate AIS-L downstream within 125 μ s of detecting an LOS or LOF defect on the incoming signal or the failure of LTE supporting provisioned line origination functions. The AIS-L shall be generated as an OC-N or STS-N electrical signal that contains valid Section overhead and a scrambled all-ones pattern for the remainder of the signal.

Page 6-43

R6-164 [513v2] STE shall deactivate AIS-L within 125 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 125 μ s of clearing the failure or determining that standby equipment has been switched in.

Page 6-44

R6-165 [514] LTE shall detect an AIS-L defect on the incoming signal when bits 6, 7, and 8 of the K2 byte contain the '111' pattern in five consecutive frames.

Page 6-44

R6-166 [515] LTE shall terminate the AIS-L defect on the incoming signal when bits 6, 7, and 8 of the K2 byte have any pattern other than '111' in five consecutive frames.

Page 6-44

R6-167 [516] An NE shall declare an AIS-L failure if an AIS-L defect persists for 2.5 (± 0.5) seconds. Upon declaring an AIS-L failure, the NE shall set

an AIS-L failure indication for the line and (if AIS-L is a Reported failure for the line) send a message to an OS unless the condition in **R6-285 [626v2]** applies.

Page 6-44

- R6-168** [517] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare AIS-L failures. Upon declaring an AIS-L failure, the NE shall perform the actions listed in **R6-167 [516]**.

Page 6-44

- R6-169** [518] An NE shall clear the AIS-L failure when the AIS-L defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-L failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-44

- R6-170** [519v3] LTE shall generate AIS-P downstream for the affected STS paths within 125 μ s of detecting an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) or (if the STS pointer is processed) an LOP-P-defect on the incoming signal, or the failure of STS PTE supporting provisioned path origination functions. The AIS-P shall be generated as all-ones in the H1, H2 and H3 bytes, and the entire STS SPE.

Page 6-45

- R6-171** [521v2] LTE shall deactivate the AIS-P within 125 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 125 μ s of clearing the failure or determining that standby equipment has been switched in. LTE that performs STS pointer processing shall deactivate AIS-P by constructing a correct STS pointer with a set NDF, followed by normal pointer operations, as well as ceasing to insert the all-ones pattern in the STS SPE. LTE that does not perform STS pointer processing shall deactivate AIS-P by ceasing the insertion of all-ones in the H1, H2 and H3 bytes, and the STS SPE.

Page 6-45

- R6-172** [523] STS PTE shall detect an AIS-P defect when the H1 and H2 bytes for an STS path contain an all-ones pattern in three consecutive frames. For an STS-Nc path, only the H1 and H2 bytes of the first STS-1 need to be observed.

Page 6-45

R6-173 [524] STS PTE shall terminate an AIS-P defect when the H1 and H2 bytes for the STS path contain a valid STS Pointer with a set NDF, or when they contain valid, identical STS Pointers with normal NDFs for three consecutive frames. For an STS-Nc path, the concatenation indicators must also be valid.

Page 6-46

O6-174 [1079] STS PTE should terminate an AIS-P defect when it detects an LOP-P defect.

Page 6-46

R6-175 [525] An NE shall declare an AIS-P failure if an AIS-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an AIS-P failure, the NE shall set an AIS-P failure indication for that path and (if AIS-P is a Reported failure for the path) send a message to an OS unless the condition in **R6-285 [626v2]** applies.

Page 6-46

R6-176 [526] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare AIS-P failures. Upon declaring an AIS-P failure, the NE shall perform the actions listed in **R6-175 [525]**.

Page 6-46

R6-177 [527] An NE shall clear an AIS-P failure when the AIS-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

Page 6-46

R6-178 [528v4] STS PTE shall generate AIS-V downstream for the affected VT paths within 500 μ s of detecting an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-P defect, an UNEQ-P defect, a TIM-P defect (if activated, also see GR-253-ILR Issue ID 253-139), a PLM-P defect, or (if the VT pointer is processed) an LOP-V defect on the incoming signal, or the failure of VT PTE supporting provisioned path origination functions. The AIS-V shall be generated as an all-ones code in the entire VT, including the V1 through V4 bytes.

Page 6-47

- R6-179** [529] VT PTE shall generate AIS-V within 500 μ s of detecting a DS1 LOS, OOF, or AIS defect on an incoming DS1 that it byte-synchronously maps into a single VT1.5.
Page 6-47
- R6-180** [531v2] STS PTE shall deactivate AIS-V within 500 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 500 μ s of clearing the failure or determining that standby equipment has been switched in. STS PTE that performs VT pointer processing shall deactivate AIS-V by constructing a correct VT pointer with valid VT size and a set NDF, followed by normal pointer operations, as well as ceasing to insert the all-ones pattern in the rest of the VT. STS PTE that does not performing VT pointer processing shall deactivate AIS-V by ceasing the insertion of the all-ones pattern in the entire VT.
Page 6-47
- R6-181** [533] VT PTE shall deactivate AIS-V (as described in **R6-180 [531v2]**) within 500 μ s of terminating the defect on the incoming DS1 that caused it to be generated.
Page 6-47
- R6-182** [534] VT PTE shall detect an AIS-V defect for the VT path upon receiving an all-ones pattern in the V1 and V2 bytes in three consecutive superframes.
Page 6-47
- R6-183** [535] VT PTE shall terminate an AIS-V defect upon receiving a valid VT Pointer with valid VT size and a set NDF, or upon receiving three consecutive superframes containing valid, identical VT Pointers with a valid VT size and normal NDFs.
Page 6-48
- O6-184** [1080] VT PTE should terminate an AIS-V defect when it detects an LOP-V defect.
Page 6-48
- R6-185** [536] An NE shall declare an AIS-V failure if an AIS-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the AIS-V failure, the NE shall set an AIS-V failure indication for the path and (if AIS-V is a Reported failure for the path) send a message to an OS unless the condition in **R6-285 [626v2]** applies.
Page 6-48

-
- R6-186** [537] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-56 [419]** to declare AIS-V failures. Upon declaring an AIS-V failure, the NE shall perform the actions listed in **R6-185 [536]**.
Page 6-48
- R6-187** [538] An NE shall clear an AIS-V failure when the AIS-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-V failure indication and send a clear message to an OS (if the failure was reported to the OS).
Page 6-49
- R6-188** [539v2] STS or VT PTE shall generate DS1, DS1C, DS2, or DS3 AIS downstream within 125 μ s of the detection of certain defects, as shown in Figures 6-5 through 6-10.
Page 6-49
- R6-189** [951] VT PTE with DS0 rearrangement capabilities shall generate DS0 AIS downstream within 3 ms of the detection of certain defects (unless it is provisioned to apply a service-specific trunk conditioning code, see Section 6.2.1.6), as shown in Figures 6-11 and 6-12.
Page 6-49
- R6-190** [540] STS or VT PTE shall remove a downstream DS1, DS1C, DS2, or DS3 AIS within 125 μ s of terminating the defect that caused it to be sent.
Page 6-49
- R6-191** [541v2] If a defect that causes VT PTE to insert DS0 AIS is terminated before a failure is declared, then the VT PTE shall remove the DS0 AIS within 3 ms of terminating that defect. If a failure was declared, then the VT PTE shall remove the DS0 AIS within 3 ms of clearing the failure.
Page 6-49
- R6-192** [542] A SONET NE shall monitor for DS_n AIS and declare DS_n AIS failures on DS_n paths that it terminates.
Page 6-50
- CR6-193** [544v2] A SONET NE with DS_n interfaces may be required to monitor for DS_n AIS (as if it were terminating the DS_n path) on incoming DS_n signals where the DS_n path is not terminated.
Page 6-50
-

-
- CR6-194** [952] A SONET NE with DS_n interfaces may be required to monitor for DS_n AIS (as if it were terminating the DS_n path) on DS_n signals that are asynchronously demultiplexed from STS or VT SPEs (i.e., on outgoing DS_n signals).
Page 6-50
- O6-195** [953] If the DS_n path is not terminated, then the capability to monitor for DS_n AIS should be provided independently of any options to provide non-clear-channel transport of DS_n signals using the asynchronous DS_n mappings (see Section 6.2.1.1.2).
Page 6-50
- R6-196** [547] A SONET NE that detects DS0 AIS shall detect a DS0 AIS defect if it receives two consecutive sets of ABCD signaling bits set to the code '0010' (i.e., the '0010' code persisting for 6 ms).
Page 6-50
- R6-197** [548] A SONET NE shall terminate a DS0 AIS defect if it receives four consecutive sets of ABCD signaling bits set to any code other than the '0010' code (i.e., any code other than '0010' persisting for 12 ms).
Page 6-51
- R6-198** [954] A SONET NE that is provisioned to insert a service-specific trunk conditioning code on a particular DS0 path shall declare a DS0 AIS failure when a DS0 AIS defect persists for 3.25 (± 0.25) seconds.
Page 6-51
- R6-199** [955] A SONET NE that is provisioned to insert a service-specific trunk conditioning code on a particular DS0 path shall clear a DS0 AIS failure when the DS0 AIS defect is terminated.
Page 6-51
- R6-200** [549v2] LTE shall generate RDI-L within 125 μ s of detecting an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2 and Figures 6-4 through 6-13). The LTE shall generate RDI-L by inserting the code '110' in bits 6, 7, and 8 of the K2 byte.
Page 6-51
- R6-201** [550v2] If bits 6 through 8 of the K2 byte are not used for other purposes (e.g., the linear APS mode indication), the LTE shall deactivate RDI-L by inserting the code '000' in bits 6, 7, and 8 of the K2 byte within 125 μ s of terminating the defect that caused it to be sent (assuming it has been sent for any minimum RDI-L assertion time supported by the NE, see below).
Page 6-52
-

-
- R6-202** [551v2] If bits 6 through 8 of the K2 byte are used for other purposes, LTE shall deactivate RDI-L by inserting an "appropriate code" (see below) in those bits within 125 μ s of terminating the defect that caused it to be sent (assuming it has been sent for the minimum assertion time).
Page 6-52
- O6-203** [956] When LTE generates RDI-L, it should generate it for at least 20 frames.
Page 6-52
- R6-204** [552v2] LTE shall detect an RDI-L defect when bits 6, 7, and 8 of the K2 byte contain the '110' pattern in five to ten consecutive frames.
Page 6-53
- R6-205** [553v2] LTE shall terminate the RDI-L defect when bits 6, 7, and 8 of the K2 byte contain any pattern other than the code '110' in five to ten consecutive frames.
Page 6-53
- R6-206** [554] An NE shall declare an RFI-L failure when an RDI-L defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-L failure, the NE shall set an RFI-L failure indication for the line and (if RFI-L is a Reported failure, see Section 6.2.1.8) send a message to an OS.
Page 6-53
- R6-207** [555] An NE shall clear the RFI-L failure when the RDI-L defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the RFI-L failure indication and send a clear message to an OS (if the failure was reported to the OS).
Page 6-53
- O6-208** [957v2] An NE should support ERDI-P generation and detection.
Page 6-54
- R6-209** [556v2] STS PTE shall generate an appropriate RDI-P signal, as specified in Table 6-4, within 100 ms of detecting a listed defect. (Also see Figures 6-5 through 6-13.)
Page 6-54
- O6-210** [557v2] STS PTE should generate an appropriate RDI-P signal as specified in Table 6-4 within 125 μ s of detecting a listed defect.
Page 6-54
-

-
- R6-211** [958] If ERDI-P is supported and the STS PTE has detected two or more of the listed defects, it shall generate the higher priority ERDI-P code based on the priorities shown in Table 6-4. Page 6-54
- R6-212** [558] When STS PTE generates a particular type of RDI-P signal, it shall generate it for at least 10 frames. Page 6-54
- O6-213** [959] When STS PTE generates a particular type of RDI-P signal, it should generate it for at least 20 frames. Page 6-54
- R6-214** [559v2] STS PTE shall deactivate (or change, as appropriate) an RDI-P signal within 100 ms of terminating the defect that caused it to be generated. Page 6-55
- O6-215** [560v2] STS PTE should deactivate the RDI-P signal within 125 μ s of terminating the defect that caused it to be sent (assuming that the signal has been sent for the minimum assertion time supported by the NE, see **R6-212** [558] and **O6-213** [959]). Page 6-55
- R6-216** [561] If **O6-210** [557v2] and **O6-215** [560v2] are not both met, then the delay time to generate an RDI-P signal (i.e., the time between detection of the defect and generation of the RDI-P signal) shall be within 500 μ s of the delay time to deactivate the RDI-P signal (i.e., the time between termination of the defect and deactivation of the RDI-P signal). Page 6-55
- R6-217** [960] STS PTE that does not support ERDI-P shall detect a one-bit RDI-P defect when a '1' is received in bit 5 of G1 for ten consecutive frames. Page 6-55
- R6-218** [562v2] STS PTE that supports ERDI-P shall detect an RDI-P defect when one of the "RDI-P defect" codes shown in Table 6-4 (one-bit or enhanced) is received for five to ten consecutive frames. Page 6-55
-

-
- R6-219** [961] STS PTE that does not support ERDI-P shall terminate the one-bit RDI-P defect when a '0' is received in bit 5 of G1 for ten consecutive frames.
Page 6-55
- R6-220** [564v2] STS PTE that supports ERDI-P shall terminate a particular type of RDI-P defect (one-bit or enhanced) when a code other than the code corresponding to that defect is received for five to ten consecutive frames.
Page 6-55
- R6-221** [566v2] An NE shall declare the corresponding one-bit RFI-P or ERFI-P failure when a particular type of RDI-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-P failure, the NE shall set an RFI-P failure indication for the STS path and (if RFI-P is a Reported failure, see Section 6.2.1.8) send a message to an OS unless the condition in **R6-286 [629v3]** applies.
Page 6-55
- R6-222** [962v2] If ERDI-P is supported, the RFI-P failure indication shall indicate if the failure was derived from an incoming Server, Connectivity or Payload ERDI-P defect, or a one-bit RDI-P defect (see Table 6-4).
Page 6-56
- R6-223** [567v2] An NE shall clear the corresponding RFI-P failure when the particular type of RDI-P defect that caused it to be declared is absent for 10 (± 0.5) seconds. Upon clearing the RFI-P failure, the NE shall clear the RFI-P failure indication and send a clear message to the OS (if the failure was reported to the OS).
Page 6-56
- O6-224** [963v2] An NE should support enhanced RDI-V generation and detection.
Page 6-57
- R6-225** [568v2] VT PTE shall generate an appropriate RDI-V signal, as specified in Table 6-5, within 100 ms of detecting a listed defect. (Also see Figures 6-7 through 6-13.)
Page 6-57
- O6-226** [569v2] VT PTE should generate an appropriate RDI-V signal as specified in Table 6-5 within 500 μ s of detecting a listed defect.
Page 6-57
-

-
- R6-227** [964] If ERDI-V is supported and the VT PTE has detected two or more of the listed defects, it shall generate the higher priority ERDI-V code based on the priorities shown in Table 6-5.
Page 6-57
- R6-228** [570] When VT PTE generates a particular type of RDI-V signal, it shall generate that signal for at least 10 superframes.
Page 6-58
- O6-229** [965] When VT PTE generates a particular type of RDI-V signal, it should generate it for at least 20 superframes.
Page 6-58
- R6-230** [571v2] VT PTE shall deactivate (or change, as appropriate) an RDI-V signal within 100 ms of terminating the defect that caused it to be sent.
Page 6-58
- O6-231** [572v2] VT PTE should deactivate the RDI-V signal within 500 μ s of terminating the defect that caused it to be sent (assuming that the signal has been sent for the minimum assertion time supported by the NE, see **R6-228** [570] and **O6-229** [965]).
Page 6-58
- R6-232** [573] If **O6-226** [569v2] and **O6-231** [572v2] are not both met, then the delay time to generate the RDI-V signal (i.e., the time between detection of the defect and generation of the RDI-V signal) shall be within 2 ms of the delay time to deactivate the RDI-V signal (i.e., the time between termination of the defect and deactivation of the RDI-V signal).
Page 6-58
- R6-233** [966] VT PTE that does not support ERDI-V shall detect a one-bit RDI-V defect when a '1' is received in bit 8 of V5 for ten consecutive superframes.
Page 6-58
- R6-234** [574v2] VT PTE that supports ERDI-V shall detect an RDI-V defect when one of the "RDI-V defect" codes shown in Table 6-5 (one-bit or enhanced) is received for five to ten consecutive VT superframes.
Page 6-59
- R6-235** [967] VT PTE that does not support ERDI-V shall terminate the one-bit RDI-V defect when a '0' is received in bit 8 of V5 for ten consecutive superframes.
Page 6-59
-

R6-236 [576v2] VT PTE that supports ERDI-V shall terminate a particular type of RDI-V defect (one-bit or enhanced) when a code other than the code corresponding to that defect is received for five to ten consecutive superframes.

Page 6-59

R6-237 [578v2] An NE that is not using the byte-synchronous DS1 mapping for a particular VT path shall declare the corresponding one-bit RFI-V or RFI-V failure when a particular type of RDI-V defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-V failure, the NE shall set an RFI-V failure indication for the VT path and (if RFI-V is a Reported failure, see Section 6.2.1.8) send a message to an OS unless the condition in **R6-286 [629v3]** applies.

Page 6-59

R6-238 [968v2] If ERDI-V is supported, the RFI-V failure indication shall indicate if the failure was derived from an incoming Server, Connectivity or Payload ERDI-V defect, or a one-bit RDI-V defect (see Table 6-5).

Page 6-59

R6-239 [579v2] An NE that is not using the byte-synchronous DS1 mapping for a particular VT path shall clear the corresponding RFI-V failure when the particular type of RDI-V defect that caused it to be declared is absent for 10 (± 0.5) seconds. Upon clearing the RFI-V failure, the NE shall clear the RFI-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

Page 6-59

R6-240 [580v3] VT PTE that is using the byte-synchronous DS1 mapping for a particular path shall generate an RFI-V signal by setting bit 4 of the V5 byte to '1' within 500 μ s of declaring an AIS-V failure (or a lower-layer, traffic-related, near-end failure, see Section 6.2.1.8.2), an LOP-V failure, an UNEQ-V failure, or a PLM-V failure (as shown in Figures 6-9 through 6-13).

Page 6-60

R6-241 [969] VT PTE that is byte-synchronously mapping a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall generate an RFI-V signal by setting bit 4 of the V5 byte to '1' within 500 μ s of declaring a DS1 RAI failure for an incoming DS1 (as shown in Figure 6-10).

Page 6-60

-
- R6-242** [581v2] VT PTE that is using the byte-synchronous DS1 mapping shall deactivate the RFI-V signal by setting bit 4 of V5 to '0' within 500 μ s of clearing the failure that caused the RFI-V signal to be sent.
Page 6-61
- R6-243** [582v2] An NE that is using the byte-synchronous DS1 mapping for a particular VT path shall declare an RFI-V failure upon receiving a '1' in bit 4 of V5 for 10 consecutive superframes (i.e., for 5 ms). Upon declaring the RFI-V failure, the NE shall set an RFI-V failure indication for the VT path and (if RFI-V is a Reported failure, see Section 6.2.1.8) send a message to an OS unless the condition in **R6-286** [629v3] applies.
Page 6-61
- R6-244** [583v2] An NE shall clear the RFI-V failure within 50 ms of receiving a '0' in bit 4 of V5 for 10 consecutive superframes. Upon clearing the RFI-V failure, the NE shall clear the RFI-V failure indication and send a clear message to the OS (if the failure was reported to the OS).
Page 6-61
- CR6-245** [970] A SONET NE that terminates an M23 application DS3 path may be required to generate DS3 RDI upstream upon detecting certain defects, as shown in Figure 6-8.
Page 6-62
- R6-246** [971] A SONET NE that terminates a C-bit parity application DS3 path shall generate DS3 RDI upstream upon detecting certain defects, as shown in Figure 6-8.
Page 6-62
- R6-247** [972] A SONET NE shall remove a DS3 RDI upon terminating the defect that caused it to be generated.
Page 6-62
- R6-248** [584v2] A SONET NE that terminates a DS_n path other than an M23 application DS3 path shall generate DS_n RAI upstream (unless it is provisioned to apply a service-specific trunk conditioning code upstream) upon declaring certain failures, as shown in Figures 6-8, 6-12 and 6-13.
Page 6-62
- R6-249** [586v2] VT PTE provisioned to byte-synchronously map a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall generate DS1 RAI downstream upon declaring an RFI-V failure as a result of receiving an RFI-V signal, as shown in Figure 6-9.
Page 6-62
-

-
- R6-250** [973] A SONET NE that provides access to and processing of individual DS0 channels shall generate DS0 RAI upstream upon declaring certain failures that cause trunk conditioning to be applied downstream (unless it is provisioned to also apply a service-specific trunk conditioning code upstream), as shown in Figures 6-11 and 6-12.
Page 6-62
- R6-251** [974] A SONET NE that provides access to and processing of individual DS0 channels shall generate DS0 RAI downstream when it declares a DS1 RAI failure (unless it is provisioned to apply a service-specific trunk conditioning code), as shown in Figure 6-12.
Page 6-62
- R6-252** [975] A SONET NE shall remove a DS_n RAI upon clearing the failure that caused it to be generated.
Page 6-62
- CR6-253** [976] A SONET NE that terminates M23 application DS3 paths may be required to detect DS3 RDI for those paths.
Page 6-63
- R6-254** [977] A SONET NE that terminates C-bit parity application DS3 paths shall detect DS3 RDI for those paths.
Page 6-63
- R6-255** [978] A SONET NE that terminates DS_n paths shall detect DS_n RAI signals (if defined) for those paths.
Page 6-63
- R6-256** [585v2] VT PTE that byte-synchronously maps a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall detect a DS1 RAI signal received at the DS1 interface (and insert RFI-V downstream as shown in Figure 6-10).
Page 6-63
- R6-257** [979] A SONET NE that provides access to and processing of individual DS0 channels, and that is provisioned to apply trunk conditioning in a particular direction shall detect DS0 RAI signals in that direction (and apply the trunk conditioning code).
Page 6-63
- R6-258** [587v2] A SONET NE that detects DS0 RAI signals shall declare a DS0 RAI failure if it receives two consecutive sets of ABCD signaling bits set to the code '0111' (i.e., the code '0111' persisting for 6 ms).
Page 6-63
-

-
- R6-259** [588v2] A SONET NE shall clear a DS0 RAI failure if it receives four consecutive sets of ABCD signaling bits set to any code other than the '0111' code (i.e., any code other than '0111' persisting for 12 ms).
Page 6-63
- CR6-260** [591v2] STS PTE may be required to support PDI-P signal generation.
Page 6-64
- CR6-261** [589v2] STS PTE that supports PDI-P generation may be required to be provisionable as to whether it sends PDI-P signals.
Page 6-64
- R6-262** [980] If STS PTE that supports PDI-P generation and a VT-structured STS SPE does not process VT pointers, it shall (non-intrusively) detect and terminate LOP-V defects as if it were processing the those pointers (i.e., according to the criteria in Section 6.2.1.1.3).
Page 6-64
- R6-263** [981] STS PTE that supports PDI-P generation and a VT-structured STS SPE shall (non-intrusively) detect and terminate AIS-V defects as if it were VT PTE (i.e., according to the criteria in Section 6.2.1.2.3).
Page 6-64
- R6-264** [982] STS PTE that supports PDI-P generation and the asynchronous DS3 mapping shall (non-intrusively) detect and terminate DS3 AIS as if it were terminating the DS3 path (i.e., according to the criteria in Section 6.2.1.2.4).
Page 6-64
- CR6-265** [593v2] STS PTE that supports PDI-P generation and the asynchronous DS3 mapping may be required to be provisionable to (non-intrusively) detect and terminate DS3 OOF as if it were terminating the DS3 path (i.e., according to the criteria in Section 6.2.1.1.2).
Page 6-64
- R6-266** [983] As a default STS PTE shall not monitor for DS3 OOF for the purposes of generating PDI-P.
Page 6-65
- R6-267** [592v2] STS PTE that supports PDI-P generation shall generate (or change, as appropriate) the PDI-P signal within 100 ms of detecting an LOP-V, AIS-V, DS3 AIS, DS3 LOS, or (if so provisioned) DS3 OOF defect on any VT or DS3 payload that it embeds into the STS SPE that it
-

is originating. The PDI-P signal shall be generated by inserting the code indicating the number of defective payloads as specified in Table 3-3.

Page 6-65

- R6-268** [595v2] STS PTE that supports PDI-P generation shall deactivate (or change, as appropriate) the PDI-P signal within 100 ms of terminating a defect on one or more of its payloads.

Page 6-65

- CR6-269** [984] A SONET NE may be required to support the detection of PDI-P signals.

Page 6-65

- CR6-270** [985] A SONET NE that supports PDI-P detection may be required to be provisionable (on a per-STs path basis) as to whether it detects PDI-P.

Page 6-65

- R6-271** [596v2] A SONET NE that supports the detection of PDI-P shall detect a PDI-P defect (or a change in the PDI-P defect, as appropriate) within 10 ms of the onset of at least five consecutive samples (which might not be consecutive frames) of STS Signal Labels (C2 bytes) containing a new PDI-P code.

Page 6-65

- O6-272** [597v2] A SONET NE that supports the detection of PDI-P should detect a PDI-P defect (or a change in the PDI-P defect, as appropriate) immediately upon receipt of five consecutive frames of STS Signal Labels containing a new PDI-P code.

Page 6-65

- R6-273** [598] A SONET NE that supports the detection of PDI-P shall terminate a PDI-P defect within 10 ms of the onset of at least five consecutive samples (which might not be consecutive frames) of STS Signal Labels that do not contain a PDI-P code.

Page 6-65

- O6-274** [599] A SONET NE that supports the detection of PDI-P should terminate a PDI-P defect immediately upon receipt of five consecutive frames of STS Signal Labels that do not contain a PDI-P code.

Page 6-65

- R6-275** [610v3] An NE that provides access to and processing of individual DS0 channels shall set a red alarm when it declares an LOS, LOF, LOP-P,

UNEQ-P, TIM-P (if activated), PLM-P, LOP-V, UNEQ-V, or PLM-V failure on the incoming signal.

Page 6-67

R6-276 [612] An NE shall clear a red alarm when the failure that caused it to be set has cleared.

Page 6-67

R6-277 [986] An NE that provides access to and processing of individual DS0 channels shall allow the user to provision it to apply trunk conditioning codes (for use in place of downstream DS0 AIS, downstream DS0 RAI, or upstream DS0 RAI, as applicable) on a per-DS0 basis. If the DS0 is not terminated, then the user shall be able to provision separate codes in each direction.

Page 6-67

R6-278 [609v2] If it is provisioned to apply a trunk conditioning code on a particular DS0, the SONET NE shall freeze the DS0 signaling state upon detecting a defect that would otherwise cause DS0 AIS to be generated or passed downstream as shown in Figures 6-11 and 6-12.

Page 6-67

R6-279 [611v2] If a defect that causes a SONET NE to freeze DS0 signaling states persists so that the NE declares the associated failure, the NE shall apply the provisioned downstream trunk conditioning codes.

Page 6-67

R6-280 [987] If a defect that causes a SONET NE to freeze DS0 signaling states persists so that it declares a failure, the NE shall apply the provisioned upstream trunk conditioning codes (if any) instead of DS0 RAI.

Page 6-68

R6-281 [617v2] If it is provisioned to apply a trunk conditioning code on a particular DS0, the SONET NE shall apply the provisioned code downstream upon declaring an RFI-V, DS1 RAI, or DS0 RAI failure as a result of an incoming RFI-V, DS1 RAI, or DS0 RAI signal (as shown in Figures 6-11 and 6-12).

Page 6-68

R6-282 [618v2] A SONET NE shall remove trunk conditioning upon clearing the failure that had caused it to be applied.

Page 6-68

-
- R6-283** [621] AIS and RFI failures shall have a default setting of Not Reported.
Page 6-84
- R6-284** [622v2] The Far-End Protection Line failure shall have a default setting of a MN alarm.
Page 6-84
- R6-285** [626v2] A SONET NE shall not autonomously report a near-end failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE, per the hierarchy in Table 6-6. In addition, the SONET NE shall not autonomously report a near-end failure declared for equipment (e.g., STS PTE) that has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8).
Page 6-85
- R6-286** [629v3] An NE that is set to report RFI failures shall not autonomously report an RFI failure that is apparently caused by the same root-cause incoming signal problem or maintenance signal at the far-end that caused the NE to concurrently declare (and report) a higher-priority RFI failure, per the hierarchy in Table 6-7. In addition, the SONET NE shall not autonomously report a far-end failure declared for equipment (e.g., STS PTE) that has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8).
Page 6-86
- R6-287** [988] The declaration of a "new" failure by a SONET NE shall not automatically cause the NE to clear any previously declared, independent failures.
Page 6-88
- R6-288** [625] A SONET NE shall provide the capability to report, on-demand to the user, the current failure indications (i.e., the current condition of the NE).
Page 6-89
- R6-289** [627v2] A SONET NE shall not report a failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE (per the hierarchy in Table 6-6) in response to a request to report all failures at the NE.
Page 6-89
- R6-290** [628v2] A SONET NE shall report all failures, including each failure that is the result of the same root-cause incoming signal problem or
-

maintenance signal as another failure reported by the NE, in response to a request to report all failures at a given SONET layer, or for a given entity.

Page 6-89

R6-291 [620] SONET equipment shall provide the capability of setting any AIS (including DS_n AIS), RFI (including DS_n RAI), and Far-End Protection Line failure as either Reported or Not Reported, and if Reported, as either alarmed or Not Alarmed. The settings shall be provisionable on a per-layer, per-entity basis (e.g., for the line layer, the settings shall be provisionable on a per-line basis).

Page 6-90

R6-292 [623] A SONET NE shall provide the capability of reporting (on demand) all software settable attributes.

Page 6-90

R6-293 [632] A SONET NE shall individually clear (and send a clear message to the OS) any failure that is individually reported to an OS.

Page 6-90

R6-294 [695v2] An NE that is non-intrusively monitoring a signal shall declare and clear failures for that signal as if it were terminating the signal. Upon declaring or clearing a failure, the NE shall set or clear the appropriate failure indication for that signal.

Page 6-91

O6-295 [703v3] As a default, an NE that is non-intrusively monitoring a signal should not autonomously report to an OS the declaration or clearing of a failure on that signal (i.e., the default or "fixed" setting for the failures should be Not Reported).

Page 6-91

R6-296 [633] Except as specifically noted, SONET NEs shall meet the general PM requirements in GR-820-CORE.

Page 6-92

R6-297 [634v2] For each PM parameter accumulated for a Physical, Section, Line, STS Path or VT Path layer entity, a SONET NE shall provide one current 15-minute, one current day, one previous 15-minute, one previous day, and 31 recent 15-minute accumulation and storage registers.

Page 6-98

- R6-298** [639v2] The size of the PM parameter accumulation registers provided by a SONET NE shall be greater than or equal to the minimum accumulation register sizes shown in Table 6-8.
Page 6-98
- R6-299** [636] A SONET NE shall allow the user to initialize all current 15-minute or current-day registers to zero at any time, on an individual entity (e.g., STS Path) basis, per direction (e.g., near-end).
Page 6-98
- R6-300** [637v2] At the end of each 15-minute period, the current 15-minute register, the previous 15-minute register, and the 31 recent 15-minute registers shall behave as a push-down stack. The current 15-minute registers shall then automatically be initialized to zero.
Page 6-98
- R6-301** [989] At the end of each day, the contents of each current-day register shall be copied to the corresponding previous day register. The current day registers shall then automatically be initialized to zero.
Page 6-98
- R6-302** [635v3] A SONET NE shall provide, as a minimum, an invalid-data flag associated with each monitored entity's near-end parameters, and another invalid-data flag associated with its far-end parameters (if applicable).
Page 6-99
- R6-303** [990] Invalid-data flags shall be moved with the data to which they apply.
Page 6-99
- R6-304** [638v3] The following apply regarding threshold registers:
- • 15-minute and 1-day threshold registers shall be provided for the SONET PM parameters that require thresholding (see Sections 6.2.2.4 through 6.2.2.7)
 - ... • The size of the threshold registers shall be greater than or equal to the minimum threshold register sizes shown in Table 6-8
 - ... • The value in each threshold register (with the possible exception of certain Physical layer threshold registers, see Section 6.2.2.2.2) shall be provisionable
 - ... • Default values for all threshold registers shall be provided and documented by the NE supplier

- ...
 - Where equivalent near-end and far-end PM parameters are defined, the default threshold value shall be the same for the near-end and the far-end parameters
- ...
 - For PM parameters where default threshold values are shown in Table 6-8, the default values provided by the NE shall be equal to those values.

Page 6-99

- R6-305** [1081] Unless the applicable equipment (e.g., the optical transmitter) has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8.2), an out-of-range alert shall be sent to an OS when the new "snapshot" value in a current Physical layer 15-minute or current day register is not in the acceptable range for that parameter [i.e., when it is less than or equal to the lower threshold (if defined), or is greater than or equal to the upper threshold].

Page 6-99

- R6-306** [991v2] Unless the applicable equipment (e.g., the LTE) has been provisioned to a service state in which autonomous reporting is inhibited (see Section 6.2.1.8.2), a TCA shall be sent to an OS when the value in a current non-Physical layer 15-minute or current day register reaches or exceeds the corresponding threshold and it is not possible for that value to be adjusted to less than the threshold based on entry into or exit from unavailable time.

Page 6-100

- R6-307** [992v2] The following apply regarding the adjustment of PM data in previous period registers (due to entry into or exit from unavailable time) and the generation of TCAs.

- ...
 - An NE shall either provide the capability to adjust the data stored in its most recent previous period registers based on entry into or exit from unavailable time during the first 10 seconds of a new current period, or it shall delay updating the values in its current period registers (and the generation of TCAs and the moving of data between registers) for up to 10 seconds.
- ...
 - If the capability to adjust the data stored in previous period registers is supported, a TCA shall be sent to an OS when the value in a previous 15-minute or previous day register is adjusted so that it is greater than or equal to the corresponding threshold.
- ...
 - If the capability to adjust the data stored in previous period registers is supported, a TCA shall be sent to an OS when the value in a previous

15-minute or previous day register can no longer be adjusted so that it will be less than the corresponding threshold (i.e., when the potential entry into or exit from unavailable time that was inhibiting the generation of the TCA at the end of a period does not occur).

- If the NE delays updating its registers (and the generation of TCAs and the moving of data between registers) for up to 10 seconds, it shall use that delay time to determine if anything has occurred that should affect the data about to be reflected in the current period registers.

Page 6-100

R6-308 [640v3] The SONET NE shall provide the capability for the user to retrieve the contents of any PM parameter register at any time. It shall also provide the capability to retrieve the contents of any threshold register (i.e., the provisioned threshold value).

Page 6-101

R6-309 [641] The SONET NE shall allow the user to schedule, and shall then perform periodic (and automatic) reporting of PM data for a monitored entity. The NE shall continue to send the appropriate PM data according to the schedule until instructed to stop by the user. This instruction to stop could be part of the scheduling information that started the periodic reporting, or it could be a separate request. The NE shall support both methods of stopping periodic performance reporting.

Page 6-101

R6-310 [642] The SONET NE shall support the ability for the user to retrieve periodic PM report schedule information.

Page 6-101

O6-311 [643v4] A SONET NE should support the LBC_{normal} and OPT_{normal} parameters to provide measurements of the health of each optical transmitter.

Page 6-106

O6-312 [994v3] A SONET NE should support the OPR_{normal} parameter to provide a measurement of the physical layer characteristics of the incoming signal at each optical receiver.

Page 6-106

R6-313 [1082] For each Physical layer PM parameter that it supports, the SONET NE shall measure and record the parameter value once per period. This

snapshot value shall be recorded at approximately the same time (i.e., within ± 10 seconds) after the start of each new period.

Page 6-106

- O6-314** [1083] The SONET NE should record its Physical layer PM parameter snapshot values within one minute after the start of each new period.
Page 6-106

- O6-315** [1084] A SONET NE should record a new snapshot value within one minute after the current period register for a Physical layer PM parameter is initialized by the user.
Page 6-106

- R6-316** [645v3] A SONET NE shall provide lower threshold registers for the OPT_{normal} and OPR_{normal} parameters (if supported) and shall also provide an upper threshold register for each Physical layer PM parameter that it supports.
Page 6-107

- R6-317** [656v2] A SONET NE shall be capable of accumulating the SEFS-S parameter.
Page 6-108

- R6-318** [657] A SONET NE shall perform thresholding for the SEFS-S parameter.
Page 6-108

- CR6-319** [658] A SONET NE may be required to support applications where line and section spans are not coincident.
Page 6-108

- R6-320** [659v2] A SONET NE that supports applications where line and section spans are not coincident shall be capable of accumulating CV-Ss, ES-Ss, and SES-Ss.
Page 6-108

- R6-321** [660v2] A SONET NE shall perform thresholding for the CV-S, ES-S and SES-S parameters if those parameters are supported.
Page 6-108

- O6-322** [1033] A SONET NE that supports applications where line and section spans are not coincident should provided one of the following capabilities:

-
- ...
- A user-provisionable, per-section option to accumulate the CV-S, ES-S and SES-S parameters that is independent of any user-provisionable option to accumulate the SEFS-S parameter.
- ...
- A user-provisionable, per-section option to ignore the B1 byte for the purposes of accumulating the CV-S, ES-S and SES-S parameters.
- Page 6-109
- R6-323** [661] A SONET NE providing LTE functions shall accumulate CV-Ls, ES-Ls, SES-Ls, UAS-Ls, and FC-Ls for each line.
- Page 6-112
- R6-324** [662] A SONET NE providing LTE functions shall perform thresholding for the CV-L, ES-L, SES-L, and UAS-L parameters.
- Page 6-112
- R6-325** [663] A SONET NE that supports line protection switching for a given line shall accumulate PSCs for that line.
- Page 6-112
- R6-326** [664] A SONET NE that is using revertive protection switching for a given line shall accumulate PSDs for that line.
- Page 6-112
- R6-327** [667] A SONET NE providing LTE functions shall provide the capability, on a per-line basis, to accumulate the CV-LFE, ES-LFE, SES-LFE, UAS-LFE, and FC-LFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be "not active."
- Page 6-112
- R6-328** [668] A SONET NE that is accumulating far-end line PM parameters shall perform thresholding for the CV-LFE, ES-LFE, SES-LFE, and UAS-LFE parameters.
- Page 6-112
- R6-329** [669] A SONET NE providing STS PTE functions shall accumulate CV-Ps, ES-Ps, SES-Ps, UAS-Ps, and FC-Ps for each terminated STS Path.
- Page 6-116
- R6-330** [670] A SONET NE providing STS PTE functions shall perform thresholding for the CV-P, ES-P, SES-P, and UAS-P parameters.
- Page 6-116
-

-
- CR6-331** [1085] STS PTE that processes the STS pointer from an incoming SONET signal (i.e., STS PTE that terminates a path that has not been reconciled to the local clock by upstream LTE within the same NE) may be required to support the capability to accumulate the STS PJ-related parameters defined in Section 6.2.2.5.1.
Page 6-116
- R6-332** [1086] If the accumulation of the STS PJ-related parameters is supported, the user shall be able to activate that accumulation on a per-path basis. The default setting shall be "not active".
Page 6-116
- R6-333** [1087] A SONET NE that is accumulating STS PJ-related parameters shall perform thresholding for the PPJC-PDet, NPJC-PDet, PPJC-PGen, NPJC-PGen, PJCDiff-P, PJCS-PDet and PJCS-PGen parameters.
Page 6-116
- R6-334** [672] A SONET NE providing STS PTE functions shall provide the capability, on a per STS Path basis, to accumulate the CV-PFE, ES-PFE, SES-PFE, UAS-PFE, and FC-PFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be "not active."
Page 6-116
- R6-335** [673] A SONET NE that is accumulating far-end STS Path PM parameters shall perform thresholding for the CV-PFE, ES-PFE, SES-PFE, and UAS-PFE parameters.
Page 6-117
- R6-336** [674] A SONET NE providing VT PTE functions shall accumulate CV-Vs, ES-Vs, SES-Vs, UAS-Vs, and FC-Vs for each terminated VT Path.
Page 6-120
- R6-337** [675] A SONET providing VT PTE functions shall provide thresholding for the CV-V, ES-V, SES-V, and UAS-V parameters.
Page 6-120
- CR6-338** [1088] VT PTE that processes the VT pointer from an incoming SONET signal (i.e., VT PTE that terminates a path that has not been reconciled to the local clock by upstream STS PTE within the same NE) may be required to support the capability to accumulate the VT PJ-related parameters defined in Section 6.2.2.6.1.
Page 6-120
-

-
- R6-339** [1089] If the accumulation of the VT PJ-related parameters is supported, the user shall be able to activate that accumulation on a per-path basis. The default setting shall be "not active".
Page 6-121
- R6-340** [1090] A SONET NE that is accumulating VT PJ-related parameters shall perform thresholding for the PPJC-VDet, NPJC-VDet, PPJC-VGen, NPJC-VGen, PJCDiff-V, PJCS-VDet and PJCS-VGen parameters.
Page 6-121
- R6-341** [676] A SONET NE providing VT PTE functions shall provide the capability to accumulate, on a per VT Path basis, the CV-VFE, ES-VFE, SES-VFE, UAS-VFE, and FC-VFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be "not active."
Page 6-121
- R6-342** [677] A SONET NE that is accumulating the Far-end VT Path PM parameters shall perform thresholding for the CV-VFE, ES-VFE, SES-VFE, and UAS-VFE parameters.
Page 6-121
- R6-343** [678] A SONET NE shall provide DS_n path PM for each DS_n path that it terminates.
Page 6-121
- R6-344** [679] A SONET NE shall provide DS_n line PM for each DS_n line that it terminates.
Page 6-121
- R6-345** [680] A SONET NE shall meet the accumulation and thresholding requirements in GR-820-CORE for DS_n Line and Path PM parameters.
Page 6-121
- R6-346** [681] A SONET NE shall meet the OS reporting and data retrieval requirements in GR-820-CORE for its DS_n PM parameters.
Page 6-121
- R6-347** [683v3] A SONET NE shall inhibit the accumulation of near-end SONET PM parameters according to the rules listed below (and as summarized in Tables 6-13 through 6-19).
-

- ... • The accumulation of all near-end Line layer parameters except for UASs, FCs, PSCs, and PSDs (if applicable) shall be inhibited during periods of unavailability of the monitored line.
- ... • The accumulation of all near-end STS or VT layer parameters except for UASs and FCs shall be inhibited during periods of unavailability of the monitored path.
- ... • The accumulation of near-end CVs for a particular entity shall be inhibited for all seconds that are counted as SESs for that entity.

Page 6-122

R6-348 [995v2] A SONET NE shall inhibit the accumulation of far-end SONET PM parameters according to the rules listed below (and as summarized in Tables 6-13 through 6-19).

- ... • The accumulation of all far-end parameters except for far-end UASs and far-end FCs shall be inhibited during periods of unavailability of the line or path at the far end.
- ... • The accumulation of far-end CVs for a particular entity shall be inhibited for all seconds that are counted as (far-end) SESs for that entity.
- ... • The accumulation of all far-end Line parameters shall be inhibited during seconds in which a near-end AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) is present. An invalid-data flag shall be set for the inhibited parameters.
- ... • The accumulation of all far-end STS Path parameters shall be inhibited during seconds in which a near-end AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), a near-end LOP-P defect, or a near-end UNEQ-P defect is present. An invalid-data flag shall be set for the inhibited parameters.
- ... • The accumulation of all far-end VT Path parameters shall be inhibited during seconds in which a near-end AIS-V defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), a near-end LOP-V defect, or a near-end UNEQ-V defect is present. An invalid-data flag shall be set for the inhibited parameters.

Page 6-123

R6-349 [682] A SONET NE that provides DS_n performance monitoring shall meet the requirements in GR-820-CORE for DS_n PM during troubles.

Page 6-123

-
- | **CR6-350** [685v2] A SONET NE may be required to support non-PJ-related level-1 intermediate-path PM for some or all types of level-1 paths in the NE.
Page 6-132
- | **CR6-351** [1091] A SONET NE containing STS PTE that processes VT pointers to frequency justify (to a local clock) VT Paths carried on an incoming SONET signal may be required to support the capability to accumulate VT PJ-related level-1 intermediate path PM parameters for those paths.
Page 6-132
- | **O6-352** [1092] A SONET NE containing LTE that processes STS pointers to frequency justify (to a local clock) STS Paths carried on an incoming SONET signal should support the capability to accumulate STS PJ-related level-1 intermediate path PM parameters for those paths.
Page 6-132
- CR6-353** [686] A SONET NE may be required to support level-2 intermediate-path PM for some or all level-2 path types in the NE.
Page 6-132
- | **R6-354** [687v2] A SONET NE that provides a given type of intermediate-path PM for a particular type of level-1 path shall provide the capability to monitor any of the level-1 paths of that type passing through the NE.
Page 6-132
- | **R6-355** [688v2] A SONET NE that provides intermediate-path PM for a particular type of level-2 path shall provide the capability to monitor any of the level-2 paths of that type passing through the NE.
Page 6-132
- | **R6-356** [1093] If a SONET NE supports level-1 or level-2 intermediate path PM for a particular type of path, then the number or percentage of the paths of that type that can be simultaneously monitored shall be clearly documented.
Page 6-132
- | **R6-357** [689v2] A SONET NE that provides a given type of intermediate-path PM for a given path type shall provide the user the ability to activate the feature on a per-path basis. The default shall be "not active" for all paths.
Page 6-132
- | **O6-358** [1094] A SONET NE that provides the capability to accumulate both PJ-related and non-PJ-related intermediate-path PM parameters for an
-

STS or VT Path should allow the user to independently activate and deactivate the accumulation of those two sets of parameters.

Page 6-133

R6-359 [690] A SONET NE that provides intermediate-path PM for a given path shall not alter any bits, overhead or payload, of the monitored entity in either direction of transmission.

Page 6-133

R6-360 [691v4] If a non-PJ-related intermediate-path PM feature is active for a particular path, the NE shall perform both near-end and (if defined) far-end path PM for the path signals in both directions. Except as noted below, the monitoring shall be performed as if the NE were terminating the path in each direction, using the criteria in Sections 6.2.2.5, 6.2.2.6, and 6.2.2.7 (for STS, VT, and DS_n paths).

Page 6-133

R6-361 [1095] If a PJ-related intermediate-path PM feature is active for a particular path, the NE shall monitor that path in the direction specified. Except as noted below, the monitoring shall be performed as if the NE were terminating the path, using the criteria in Sections 6.2.2.5 and 6.2.2.6 (for STS and VT paths).

Page 6-133

O6-362 [693] An NE that is capable of performing intermediate-path PM for a particular type of level-2 path should allow the user to provision whether it also performs level-1 intermediate-path PM on the container paths.

Page 6-134

R6-363 [694v3] An NE that is performing intermediate-path PM for a particular path shall (non-intrusively) detect and terminate AIS, LOP and RDI defects (and for DS_n paths, DS_n OOF) as if it were terminating that path.

Page 6-134

R6-364 [1034v2] An NE that is performing intermediate-path PM for a particular SONET path, and that supports the detection of ERDI defects for that path, shall (non-intrusively) detect and terminate UNEQ defects as if it were terminating the path.

Page 6-134

R6-365 [996] A SONET NE that is performing level-2 intermediate-path PM for a particular path shall (non-intrusively) detect and terminate AIS, LOP,

UNEQ, PLM, and RDI defects for the container path as if it were terminating that path.

Page 6-134

R6-366 [706] Access to the fiber for fiber testing (e.g., for identifying the location of a fiber break) shall be provided.

Page 6-136

O6-367 [707] To aid in the remote sectionalization of a problem to the fiber or to one of the terminals, the ability to switch the fiber to an alternate source or an alternate receiver should be provided.

Page 6-136

O6-368 [709] A SONET NE should provide local test access to individual STS-1s and STS-3s within an OC-N or STS-N electrical signal.

Page 6-136

R6-369 [710] The SONET electrical level test access shall be in accordance with the specifications defined in Section 4.4 for STS-1 and STS-3 electrical signals.

Page 6-136

R6-370 [711] The SONET electrical level test access shall provide a non-intrusive and hitless monitor mode.

Page 6-136

R6-371 [712] The SONET electrical level test access shall provide the ability to perform intrusive tests.

Page 6-136

R6-372 [713] When intrusive testing occurs, the NE shall provide the appropriate Path AIS in the non-test direction.

Page 6-136

R6-373 [714] In NEs supporting APS, the test access shall not impair the working channel on the protection line.

Page 6-137

CR6-374 [715] A SONET NE that is VT programmable may be required to provide DS1 remote test access.

Page 6-137

-
- R6-375** [716] The DS1 test access shall meet the criteria in FR-476, *OTGR Section 6: Network Maintenance: Access and Testing*, unless deviations from this are explicitly identified within NE-specific GRs, TRs, or TAs.
Page 6-137
- R6-376** [717] When TL1 interfaces are used, SONET NEs shall use the TL1 messages of GR-834-CORE, *OTGR Section 12.4: Network Maintenance: Access and Testing Messages*, for any test access functions provided.
Page 6-137
- R6-377** [718] When CMISE/OSI interfaces are used, SONET NEs shall adhere to the object model and use the CMISE service mappings of GR-1031-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools: Metallic Test Access*, for any test access functions provided.
Page 6-137
- R6-378** [719] A SONET NE shall provide diagnostic capabilities. As a minimum, diagnostics shall be provided to detect the equipment failures listed in Section 6.2.1.1.4. The diagnostic capabilities shall run routinely and on demand. When these diagnostics are run on demand, the NE shall provide the user with the results.
Page 6-138
- O6-379** [720] An NE should support additional diagnostics that provide the ability to isolate an equipment trouble to the replaceable circuit pack or module. These diagnostics should not interfere with working services and may be run routinely or on demand.
Page 6-138
- R6-380** [721] Diagnostics that interfere with service shall not run routinely unless permitted by the user.
Page 6-138
- O6-381** [722v2] An NE that does not support the OPR_{normal} PM parameter for an optical receiver should provide an on-demand diagnostic that reports the received optical power OPR (not normalized), preferably in units of dBm.
Page 6-138
- O6-382** [723v2] An NE that does not support the OPT_{normal} PM parameter for an optical transmitter should provide an on-demand diagnostic that reports the transmitted optical power OPT (not normalized), preferably in units of dBm.
Page 6-138
-

-
- O6-383** [724v2] An NE that does not support the LBC_{normal} PM parameter for an optical transmitter should provide an on-demand diagnostic that reports the laser bias current LBC (not a normalized), preferably in microamperes (μA).
Page 6-138
- O6-384** [1096] An NE that supports the OPR_{normal} PM parameter but does not meet **O6-315 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of OPR_{normal} .
Page 6-138
- O6-385** [1097] An NE that supports the OPT_{normal} PM parameter but does not meet **O6-315 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of OPT_{normal} .
Page 6-139
- O6-386** [1098] An NE that supports the LBC_{normal} PM parameter but does not meet **O6-315 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of LBC_{normal} .
Page 6-139
- R6-387** [740v2] The SONET NE shall provide the functional equivalent of the diagnostic illustrated in Figure 6-24.
Page 6-139
- R6-388** [741] An NE shall deny a request for this diagnostic if the incoming facility associated with the receiver is In_Service or Out_of_Service-autonomous, as defined in GR-1093-CORE.
Page 6-140
- R6-389** [743] While the diagnostic is being performed, the NE shall recognize that it is in a test state, as defined in GR-1093-CORE, and it shall not take action to revert to a previous state.
Page 6-140
- R6-390** [745] The NE shall exit the test state, as defined in GR-1093-CORE, when the diagnostic is deactivated.
Page 6-140
- R6-391** [744] An NE shall notify the OS when the diagnostic has been activated and when it has been deactivated.
Page 6-140
-

-
- R6-392** [746] The NE shall determine if the receiver is able to frame on the delivered signal.
Page 6-140
- R6-393** [747] The NE shall detect errors in the Section BIP-8 (B1) byte, or if the B1 byte is not processed, in the Line BIP-8s (B2).
Page 6-140
- R6-394** [748v2] The NE shall report results indicating that the equipment could not frame, that CVs were detected, or that no framing problems or CVs were detected.
Page 6-140
- R6-395** [728] STS PTE shall provide an on-demand diagnostic to detect and report the contents of the received STS Path Trace.
Page 6-141
- CR6-396** [729] LTE with STS cross-connection capabilities may be required to provide an on-demand diagnostic to detect and report the contents of the STS Path Trace in the (nonterminated) STS path designated by the user.
Page 6-141
- O6-397** [730] STS PTE should provide a diagnostic that, when activated, continuously monitors the incoming STS Path Trace.
Page 6-141
- CR6-398** [999v2] A SONET NE may be required to support a feature to allow the user to provision, for diagnostics purposes, the "expected" path trace for each STS path that it terminates.
Page 6-141
- R6-399** [1000v2] If an NE supports a feature that allows the user to provision an expected path trace for diagnostics purposes and that feature uses ASCII characters, then the following apply:
- ... • The feature shall allow the user to enter a string of up to 62 characters
 - ... • The feature shall place no restriction on the contents of the string, except that the characters shall be ASCII printable characters.
- Page 6-141
- R6-400** [731v3] A SONET NE that is continuously monitoring an incoming path trace and does not support an expected path trace feature for diagnostics
-

| purposes shall compare the incoming path trace with a previously received path trace.

Page 6-142

| **CR6-401** [1002v2] A SONET NE that is continuously monitoring an incoming path trace and supports an expected path trace feature for diagnostics purposes, but that is not provisioned with an expected path trace, may be required to compare the incoming path trace with a previously received path trace.

Page 6-142

R6-402 [1003] As a minimum, an NE that compares the incoming path trace with a previously received path trace shall be capable of performing a case-sensitive comparison of ASCII-based path traces.

Page 6-142

O6-403 [1004] An NE that compares the incoming path trace with a previously received path trace should be capable of performing that comparison on a byte-by-byte basis (i.e., independent of whether the contents are ASCII-based or contain <CR> and <LF> characters).

Page 6-142

| **R6-404** [735v3] A SONET NE that is monitoring an incoming path trace for changes or mismatches for diagnostics purposes shall suspend the STS Path Trace monitoring function when the J1 byte cannot be accessed (e.g., when an LOP-P or AIS-P defect has been detected by the STS PTE). If the NE is monitoring for changes in the incoming path trace, the contents of the path trace before the disruption shall be used as the starting value following a restart.

Page 6-142

R6-405 [736v2] A SONET NE that has suspended monitoring of an incoming path trace because the J1 byte could not be accessed shall resume monitoring when the J1 byte can again be accessed.

Page 6-142

R6-406 [732v2] A SONET NE that is monitoring for changes of the incoming path trace shall detect when a sustained change in the path trace content occurs. Upon detecting a sustained change, the NE shall send a message to an OS. The level of the message shall be Not Alarmed, and it shall include both the previously received path trace, and the new path trace (assuming they are ASCII-based).

Page 6-142

- R6-407** [1005v2] A SONET NE that is monitoring for a mismatch between the incoming path trace and an expected path trace for diagnostics purposes shall detect when a sustained mismatch occurs. Upon detecting a sustained mismatch, the NE shall set an indication for that path and send a message to an OS. The default level of the message shall be Not Alarmed, and it shall include both the expected path trace and the new path trace (assuming they are ASCII-based).
Page 6-143
- R6-408** [1006v2] A SONET NE that is monitoring the incoming path trace for diagnostics purposes and that has detected a sustained mismatch shall detect when the incoming path trace matches the expected path trace. Upon detecting a match, the NE shall clear the indication for that path and send a clear message to the OS (if the mismatch was reported to an OS).
Page 6-143
- CR6-409** [1007v2] An NE that monitors the incoming path trace for mismatches for diagnostics purposes may be required to be user-provisionable to report a detected mismatch as alarmed or Not Alarmed.
Page 6-143
- R6-410** [737] A SONET NE shall provide an on-demand diagnostic to report the value currently being received in the STS and VT Signal Labels at STS PTE and VT PTE, respectively.
Page 6-144
- R6-411** [1008] A SONET NE that supports the detection of PDI-P defects (e.g., at an STS-level path selector in a UPSR NE) shall provide an on-demand diagnostic to report the value currently being received in the STS Signal Label contained in each STS path that it is monitoring.
Page 6-144
- CR6-412** [1009] An NE that supports the generation of PDI-P signals may be required to support an on-demand diagnostic that reports the code that is currently being inserted on the specified STS Signal Label.
Page 6-144
- R6-413** [738] A SONET NE shall provide the ability to transmit, on demand, a fully corrupted BIP value (all parity check bits inverted) for the Section, Line, STS Path, or VT Path (as appropriate to the NE). The NE shall provide the user the capability to specify the approximate duration of the corrupted BIP diagnostic as some number of units of time or some number of frames (or VT superframes for VT Path BIP).
Page 6-144

-
- O6-414** [749] A SONET NE should provide a SONET facility loopback, as illustrated in Figure 6-26.
Page 6-146
- R6-415** [750] A SONET NE shall deny a SONET facility loopback request if either the impacted direction of transmission (i.e., the direction of the looped signal) or the incoming signal to be looped is In_Service or Out_of_Service-autonomous, as defined in GR-1093-CORE.
Page 6-146
- R6-416** [751v2] When a SONET facility loopback is activated, the SONET NE shall place the associated out-of-service facilities into a test state, as defined in GR-1093-CORE.
Page 6-146
- R6-417** [756] The SONET NE shall exit the test state, as defined in GR-1093-CORE, when the loopback is deactivated.
Page 6-147
- R6-418** [755] The SONET NE shall notify the OS when the loopback has been activated and when it has been deactivated.
Page 6-147
- O6-419** [757v2] A SONET NE providing DS_n line terminations should provide DS_n terminal loopback capabilities, as shown in Figure 6-27.
Page 6-147
- O6-420** [1010] SONET NE providing DS_n line terminations should provide DS_n facility loopback capabilities, as shown in Figure 6-28.
Page 6-147
- R6-421** [759] The following control functions shall be provided:
- ... 1. Reinitialize system – System reinitialization, or “hard boot,” reloads the operating system of the NE and may affect the state of memory and other resources.
 - ... 2. Restart system – System restart, or “soft boot,” reloads only an application onto the system and not the operating system. System restart should not affect the state of nonvolatile memory and other resources. Failure states, protection switching configuration, performance parameters, and other information necessary for fault isolation should not be affected.
-

- ... 3. Reestablish a failed entity – Reestablishing failed entity involves temporary techniques to restore a service when an entity has failed. For example, when a facility fails, restoration of a service may involve reconfiguring routing tables or rerouting facilities.
- ... 4. Remove an entity from service to run tests – For such tests, traffic should be switched off the entity, and subsequent indications should be suppressed.
- ... 5. Inhibit and allow alarmed and nonalarmed indications – This capability permits the user to suppress and restart messages from the NE.
- ... 6. Check status of equipment configuration – Equipment configuration status shall be retrievable. Equipment configuration status includes the indication of the active hardware entities in replicated equipment and the active synchronization source.
- ... 7. Protection switch capabilities – See the required functions identified in Section 5.3.6.
- ... 8. Manual switch from active to standby – This capability is used for any replicated hardware or software.
- ... 9. Manual switch between synchronization sources – See the required functions identified in Section 5.4.6.

Page 6-149

- R6-422** [760] A SONET NE shall notify the OS when any control function is executed.

Page 6-149

- O7-1** [761] The equipment should be designed to minimize the investment in the frame and bay-work by using the modular design concept to minimize the costs associated with installation and the ongoing operation of the system. TR-NWT-000078, *Generic Physical Design Requirements for Telecommunications Products and Equipment*, contains physical design requirements.

Page 7-2

- R7-2** [762] The supplier shall provide appropriate documentation and training. To accomplish this in a safe and cost-effective manner, criteria for supplier documentation for NEs in TR-TSY-000454, *Supplier Documentation for Network Elements*, shall be followed. Criteria for supplier documentation for outside plant cable are in GR-20-CORE.

Page 7-2

-
- R7-3** [763] The supplier shall provide training in accordance with TR-OPT-000839, *Supplier-Provided Training Generic Requirements*.
Page 7-2
- R7-4** [764] All safety labels shall be visible to craftspersons when equipment covers are in place and when they are removed.
Page 7-3
- R7-5** [765] Voltages at or above 140 V_{dc} or 50 V_{rms ac} shall be enclosed or guarded to prevent contact. Safety labels shall be conspicuous when the guards are in place and when they are removed.
Page 7-3
- R7-6** [766] The design shall allow craftspersons safe access to parts if metal tools are to be used (e.g., insulating sleeves to guide screwdrivers to recessed potentiometers when nearby parts have hazardous voltages present).
Page 7-3
- R7-7** [767] Arrangements shall be provided to discharge large capacitors (e.g., "bleeder" resistors).
Page 7-3
- R7-8** [768] All external metal parts shall be grounded.
Page 7-3
- R7-9** [771] The fiber optic system and required optical test equipment shall be registered and certified with the Department of Health, Education and Welfare Bureau of Radiological Health as specified in 21 CFR 1040.10, *Performance Standard for Laser Products*. Documentation demonstrating system certification shall be available to assist in the determination of fiber optic safety precautions required to install, operate, and maintain the system.
Page 7-3
- R7-10** [772] The equipment involved shall conform to the applicable performance requirements, labeling requirements, and informational requirements as specified in 21 CFR 1040.10.
Page 7-3
- R7-11** [773] The fiber optic cable and required optical splicing and test equipment shall be registered and certified with the Department of Health, Education and Welfare Bureau of Radiological Health as specified in 21 CFR 1040.10. Documentation demonstrating system certification shall
-

be available to assist in the determination of fiber optic safety precautions required to install, operate, and maintain a fiber optic system.

Page 7-4

- R7-12** [774] The equipment involved shall conform to the applicable performance requirements, labeling requirements, and informational requirements as specified in 21 CFR 1040.10.

Page 7-4

- O7-13** [775] It is an objective that suppliers allow Bellcore to analyze products and processes to determine their products' compliance with this document.

Page 7-5

- R8-1** [776] All SONET Gateways shall support level 1 and level 2 IS-IS routing functions as defined in ISO 10589 and in Appendix C.

Page 8-7

- R8-2** [777] A SONET Gateway shall support the IS Role of the ES-IS routing protocol over LAN and DCC interfaces as defined in ISO 9542 and Appendix C.

Page 8-7

- R8-3** [778] A SONET Gateway shall support the IS-IS Reachable Address Prefix functionality defined in ISO 10589.

Page 8-7

- O8-4** [779] A SONET Gateway should support the level 1 partition repair function as defined in ISO 10589 and in Appendix C.

Page 8-7

- R8-5** [780] A SONET Intermediate NE shall support level 1 IS-IS routing functions as defined in ISO 10589 and in Appendix C.

Page 8-8

- R8-6** [781] A SONET Intermediate NE shall support the IS role of the ES-IS protocol over LAN and DCC interfaces as defined in ISO 9542 and Appendix C.

Page 8-8

- R8-7** [782] A SONET End NE shall support the ES role of the ES-IS protocol over LAN and DCC interfaces as defined in ISO 9542 and Appendix C.

Page 8-8

- R8-8** [783] All SONET NEs shall provide an OS/NE communications path.
Page 8-9
- R8-9** [784] A SONET NE shall be capable of being equipped with a direct OS/NE interface.
Page 8-10
- O8-10** [1035] A SONET NE should support the capability to communicate with an OS via a LAN interface.
Page 8-10
- R8-11** [785] When a stand-alone MD is used in a SONET subnetwork, the language and protocol stack for the MD/NE interface shall be identical to that for the NE/NE interface via a LAN.
Page 8-10
- R8-12** [786] All SONET NEs shall support NE/NE operations communications paths.
Page 8-11
- R8-13** [787] To support message-oriented NE/NE operations communications, a SONET NE shall be capable of being equipped with LAN and Section DCC interfaces to other NEs.
Page 8-11
- R8-14** [788] Local craftsperson access by means of a WS is required for all SONET NEs.
Page 8-11
- R8-15** [789] SONET NEs shall provide an appropriate SONET Operations Communications Interface that conforms to the protocol profiles specified in Appendix C, SONET Operations Communications Protocol Profile – Lower Layers, and Appendix D, SONET Operations Communications Protocol Profile – Upper Layers.
Page 8-12
- R8-16** [790] SONET NEs shall support CMISE as the application layer protocol for Interactive Class communications.
Page 8-12
- R8-17** [791] When file-oriented applications are supported, SONET NEs shall support FTAM as the application layer protocol as specified in GR-1250-CORE.
Page 8-12
-

- CR8-18** [792] SONET NEs may, on an interim basis, support TL1 as the application layer protocol for Interactive Class communications.
Page 8-12
- CR8-19** [793v2] The SONET OS/NE X.25 interface may, on an interim basis, support the non-OSI communications architecture (TL1 over X.25) as specified in TR-TSY-000827, *OTGR Section 11.1: Generic Operations Interfaces: Non-OSI Communications Architecture*.
Page 8-12
- O8-20** [1036] When a SONET NE supports CMISE on the OS-NE or the NE-NE interface, it should also support the X.500-based Directory Services for TMN and SONET, as defined in ANSI T1.245, for the name/address translation service.
Page 8-13
- R8-21** [794v2] At OS/NE X.25 interfaces, SONET NEs shall support the Physical Layer requirements of the TP4/CLNS Protocol Case as described in GR-828-CORE.
Page 8-13
- R8-22** [795v2] The Physical layer shall support the following 10 Mb/s baseband Media Dependent interface:
- ...
 - 10BASE-T per ANSI/IEEE Std. 802.3i-1990, (Supplement to ISO/IEC 8802-3-1990/ANSI/IEEE Std. 802.3-1990) System Configurations for Multi-segment 10 Mb/s Baseband networks (Section 13) and Twisted Pair Medium Attachment Unit (MAU) and Baseband Medium, Type 10BASE-T (Section 14).
 - ... The electrical interface and connectors shall be as specified in ISO/IEC 8802-3/ANSI/IEEE 802.3.
- Page 8-13
- CR8-23** [1011] The Physical layer may also be required to support the following 10 Mb/s baseband Media Dependent interfaces:
- ...
 - a. 10BASE2, as specified in ISO/IEC 8802-3/ANSI/IEEE 802.3
 - b. The media independent Attachment Unit Interface (AUI) as specified in ISO/IEC 8802-3/ANSI/IEEE 802.3.
- Page 8-13
- R8-24** [796] The Section DCC, a 192-kb/s channel that is carried in 3 Section overhead bytes of the first STS-1 (i.e., the D1, D2, and D3 bytes) in an

STS-N signal, shall be used as the Physical layer of the message-oriented EOC. The order of transmission is bit 1 of D1 (most significant) through bit 8 of D3 (least significant). Data Link protocols shall transmit bits across this channel by placing them into the next most significant bits.

Page 8-13

- R8-25** [797] Section EOCs shall be protected in the same way as working traffic is protected. The protection switch for the EOC shall follow the protection architecture and mode of operation (e.g., if the traffic is protected bidirectionally, the Section EOC is also protected bidirectionally; if the traffic is protected unidirectionally, then the Section EOC is also protected unidirectionally).

Page 8-14

- R8-26** [798] A SONET RGTR that accesses the Section EOC shall read the K1 and K2 bytes in both directions to determine when an EOC is being carried with working traffic (i.e., to determine when an EOC is usable).

Page 8-14

- R8-27** [799v2] At an OS/NE X.25 interface, SONET NEs shall support the Data Link layer requirements of the TP4/CLNS Protocol Case as described in GR-828-CORE.

Page 8-14

- R8-28** [800] Media Access Control functionality for the LAN shall be as specified in ISO/IEC 8802-3 and ANSI/IEEE 802.3 CSMA/CD specifications.

Page 8-15

- R8-29** [801] Logical Link Control functionality for the LAN shall be as specified in ISO/IEC 8802-2/ANSI/IEEE 802.2 LLC Class 1 Type 1 service and as described in Appendix C.

Page 8-15

- R8-30** [802] The LSAP value 0111 1111, in which the leftmost bit is the least significant bit, shall be used for the LAN. This value would be represented as 'FE' (hex).

Page 8-15

- R8-31** [803] The Data Link layer protocol for the SONET Section DCC shall be based on Link Access Protocol on the D-channel (LAPD) as specified in ITU-T Recommendation Q.921, *ISDN user-network interface - Data link layer specification*, and as described in Appendix C.

Page 8-15

-
- R8-32** [804] Both the Unacknowledged Information Transfer Service (UITs) and the Acknowledged Information Transfer Service (AITS) shall be supported. AITS shall be the default mode of operation.
Page 8-15
- R8-33** [805] The SAPI value shall be preassigned, and shall be settable locally or remotely by an OS.
Page 8-15
- R8-34** [806] The Data Link layer procedures, with the exception of the TEI management procedure, shall follow the rules ITU-T Q.921 specifies.
Page 8-16
- R8-35** [807] SAPI value of 62 shall be used for SONET Section DCC applications.
Page 8-16
- R8-36** [808] The assignment of user-side/network-side roles (and, hence, the C/R bit value) shall be settable and made before initialization of the data link.
Page 8-16
- R8-37** [809] A TEI value of 0 (zero) shall be used for SONET Section DCC applications.
Page 8-16
- R8-38** [810] When TP4 is being run over CLNP, the N-SEL portion of the NSAP shall be set to an initial value of '1D' (hex).
Page 8-16
- R8-39** [811] When TARP is being run over CLNP, the N-SEL portion of the NSAP shall be set to an initial value of 'AF' (hex).
Page 8-16
- R8-40** [812] The capability to change the value of the N-SEL when the OSI stack is reinitialized shall be supported.
Page 8-16
- R8-41** [813v3] At an OS/NE X.25 interface, SONET NEs shall support the Network Layer requirements of the CL-WAN profile (CLNS2) as described in IUT-T Recommendation Q.811 (1997), *Lower Layer protocol profiles for the Q3 and X interfaces*, except that use of the ES-IS protocol shall not be supported.
Page 8-16
-

- R8-42** [1099] At an OS/NE X.25 interface, SONET NEs shall also support the X.25 Subnetwork Service and Protocol requirements found in Section 5.3.2.4 of GR-828-CORE.
Page 8-17
- R8-43** [814] The Network layer protocol for DCCs and LANs shall be CLNP (ISO 8473) as specified in Appendix C.
Page 8-17
- R8-44** [815] The Subnetwork Dependent Convergence Function (SNDCF), as specified in ISO 8473:1988/Add. 3, and the protocol identification convention described in ISO TR 9577, shall be used for the mapping of the primitives defined for the data link services into the underlying service assumed by the CLNP.
Page 8-17
- R8-45** [816] The ISO 8473 Category 3 QoS function shall be supported.
Page 8-17
- R8-46** [817] The coding of the QoS parameter for the selection of UITS/AITS in the data link shall be as follows:
- ... a. The absence of a QoS parameter shall select AITS.
 - ... b. Bits 7 and 8 set to 1 (Globally Unique QoS) and bit 1 set to 1 shall select AITS.
 - ... c. Bits 7 and 8 set to 1 (Globally Unique QoS) and bit 1 set to 0 shall select UITS.
- Page 8-17
- CR8-47** [818] Other ISO 8473 Category 3 functions may be required except where prohibited below.
Page 8-17
- R8-48** [819v2] The following service/protocol parameters shall have the values specified below:
- ... a. Error Reporting (E/R) Flag — As stated in ANSI T1.204, the setting of E/R flag is a local matter. The default value of this flag shall be zero to avoid excessive network traffic that can result during broadcast routing.
 - ... b. Partial Source Routing — Partial source routing shall not be supported because NBSIR 88-3824-1, containing OSI implementation agreements, has identified a defect with the partial

source routing option that can cause NPDUs to loop in the network until their lifetime expires.

- ... c. Inactive Subset — Implementations shall not transmit NPDUs encoded using the ISO 8473:1988 inactive subset. Received NPDUs encoded with the inactive subset shall be discarded.
- ... d. Segmentation — The non-segmenting subset shall not be used. However, implementations shall be capable of receiving and correctly processing NPDUs that do not contain the segmentation part.
- ... e. Segmentation Permitted (SP) Flag — Implementations shall not generate data NPDUs without a segmentation part (i.e., the SP flag shall be set to 1 and the segmentation part shall be included).
- ... f. Lifetime Control — The lifetime parameter shall be used as Paragraph 6.4 of ISO 8473:1988 specifies. This parameter shall have an initial default value of at least three times the network span (number of network entities) or three times the maximum transit delay (divided by 500 ms), whichever is greater, as ISO 8473 specifies. The initial default PDU Lifetime Control shall be 10 seconds.

Page 8-18

- O8-49** [820] The CLNS Congestion Notification should be used. The default value of '0' should be used when originating NPDUs.

Page 8-18

- O8-50** [821] The mandatory and optional approaches to congestion avoidance and recovery given in NIST Publication 500-202, Part 4, Section 5.1.2.5 should be used.

Page 8-18

- R8-51** [822] The destination and source addresses used for SONET applications shall be Network Service Access Point (NSAP), as specified in ISO 8348:1993. The Domain Specific Part (DSP) shall be the ISO DCC format as described in ANSI T1.204-1993 and ANSI X3.216-1992.

Page 8-18

- R8-52** [823] The System ID field shall be assigned a 6 octet IEEE address by the equipment supplier.

Page 8-20

-
- R8-53** [824] Class 4 of ISO 8073 (8073 Add. 2, TP4) shall be supported as specified in Appendix C.
Page 8-20
- R8-54** [825] TP4 implementations shall be capable of receiving and processing all possible parameters for all possible TPDU's, dependent upon the class and optional functions implemented.
Page 8-20
- R8-55** [826] If the ISO Session Layer is being run over TP4, then the TSAP-ID shall be set to an initial value of "TT" (ASCII), i.e., '5454' (hex).
Page 8-21
- R8-56** [827] The capability to change the value of the TSAP-ID (within a range of 0 to 4 octets) when the OSI stack is reinitialized shall be supported.
Page 8-21
- R8-57** [828] An unknown parameter in any received CR TPDU shall be ignored.
Page 8-21
- R8-58** [829] Known parameters with invalid lengths in a CR or CC TPDU shall be handled as a protocol error.
Page 8-21
- R8-59** [830] Known parameters with valid lengths but invalid values in a CR TPDU shall be handled as follows:
- | | | |
|-----|--------------------------------|-------------------------------------|
| ... | a. TSAP-ID: | Send a Disconnect Request (DR) TPDU |
| ... | b. TPDU size: | Ignore parameter, use default |
| ... | c. Version: | Ignore parameter, use default |
| ... | d. Checksum: | Discard CR TPDU |
| ... | e. Alternate protocol classes: | Protocol error |
- Page 8-21
- R8-60** [831] Unrecognized or inapplicable bits of the additional options parameter shall be ignored.
Page 8-21
- R8-61** [832] The ISO Session layer shall be supported as specified in Appendix D.
Page 8-21
-

-
- R8-62** [833] If the ISO Presentation Layer is being run over the ISO Session layer, then the Session Selector parameter shall be set to an initial value of "SS" (ASCII), i.e., '5353' (hex).
Page 8-21
- R8-63** [834] The capability to change the value of the Session Selector parameter when the OSI stack is reinitialized shall be supported.
Page 8-21
- O8-64** [835] An SS-user-data size of up to 65,535 octets should be supported.
Page 8-21
- R8-65** [836] The ISO Presentation layer shall be supported as specified in Appendix D.
Page 8-22
- R8-66** [837] The P-SEL shall be no greater than 4 octets in length.
Page 8-22
- R8-67** [838] When CMISE PDUs are sent, the P-SEL shall be set to an initial value of '01' (hex).
Page 8-22
- R8-68** [1037] When X.500 Directory Access Protocol (DAP) PDUs are sent from the Directory User Agent (DUA) to the Directory System Agent (DSA), the P-SEL shall be set to an initial value of '04' (hex).
Page 8-22
- R8-69** [839] When FTAM PDUs are sent, the P-SEL shall be set to an initial value of '02' (hex).
Page 8-22
- R8-70** [841] When TL1 PDUs are sent, the P-SEL shall be set to an initial value of 'AF' (hex).
Page 8-22
- R8-71** [842] The capability to change the value of the P-SEL when the OSI stack is reinitialized shall be supported.
Page 8-22
- R8-72** [843] The Application Control Service Element (ACSE) shall be supported as specified in Appendix D.
Page 8-22
-

-
- R8-73** [844] SONET NEs shall support ROSE/CMISE as specified by GR-828-CORE.
Page 8-23
- R8-74** [845] The TMN Application Context defined in CCITT M.3100, Section 10, shall be used.
Page 8-23
- R8-75** [846] The CMISE objects and service mappings for SONET that are contained in GR-1042-CORE and GR-1042-IMD, and the objects and service mappings supporting surveillance and memory administration that are contained in GR-836-CORE and GR-836-IMD shall be supported as per the requirements in those documents.
Page 8-23
- R8-76** [847] The Presentation context shall contain the TL1 abstract syntax and TL1 transfer syntax that these identifiers specify:
- ... tl1AbstractSyntax OBJECT IDENTIFIER ::= {1 3 17 104 11 2
bellcoreSONETSyntax(1)}
- ... tl1TransferSyntax OBJECT IDENTIFIER ::= {1 3 17 104 12 2
bellcoreSONETSyntax(1)}.
- Page 8-24
- R8-77** [848] The transfer syntax for TL1 messages shall be the ASCII encoding of the character string constituting each TL1 message.
Page 8-24
- R8-78** [849] The defined context set shall contain the following:
- ... ABSTRACT SYNTAX {1 3 17 104 11 2 bellcoreSONETSyntax(1)}
{ joint-iso-ccitt association-control (2), abstract-syntax (1), apdus (0), version (1) }
- ... TRANSFER SYNTAX {1 3 17 104 12 2 bellcoreSONETSyntax(1)}
{ joint-iso-ccitt asn1 (1), basic-encoding (1) }
- Page 8-24
- R8-79** [850v2] Presentation layer PDUs containing TL1 messages exchanged between communicating SONET NEs shall use the choice of "fully encoded data" for user data defined in CCITT X.226. Also, Presentation data values shall use the "octet-aligned" choice as shown below.
Page 8-24
-

- R8-80** [851] The TL1-ASE shall consist of the appropriate subset of the TL1 messages defined in GR-833-CORE, GR-199-CORE, and GR-834-CORE.
Page 8-25
- R8-81** [852] All Application context definitions for associations over which TL1 messages will be exchanged shall include ACSE and the TL1-ASE.
Page 8-25
- R8-82** [853] TL1 messages shall be exchanged using Data Presentation Protocol Data Units (TD PPDUs).
Page 8-25
- R8-83** [1038] The upper limit on TL1 message size at the application layer, specified in TR-TSY-000827 as 4096 bytes for TL1 over X.25, shall be 4096 bytes for TL1 over any protocol stack or transport mechanism.
Page 8-25
- R8-84** [854] Peer NE/NE communications shall use an association established with the Application context identifier below:
... tl1PeerComm OBJECT IDENTIFIER ::= {1 3 17 104 10 3 tl1PeerComm(1)}
Page 8-25
- R8-85** [855] The GNE shall maintain static information to map subtending NEs' TIDs to NSAP addresses.
Page 8-27
- R8-86** [856] If a GNE supports multiple RNEs/VC, then for an OS-NE message, the GNE shall:
... a. receive the TL1 message from an X.25 VC
... b. extract the TID information from the message and map the TID to the destination NE's NSAP
... c. determine the appropriate association to be used or established
... d. establish the association (if necessary)
... e. forward the TL1 message over the appropriate association to the destination NE.
Page 8-27
- R8-87** [857] If a GNE supports single RNE/VC, then the following requirements apply:

- ... a. the GNE shall maintain dynamic information to map X.25 LCNs to NSAP addresses of subtending NEs
- ... b. when an X.25 VC is established between an OS and a GNE for communications with a remote NE, the GNE shall
 - listen on that VC for a TL1 command
 - extract the TID from the command
 - map the TID and the X.25 VC LCN to the NSAP
 - establish the association with the remote NE
- ... c. for an OS-to-NE message, the GNE shall
 - receive the TL1 message from an X.25 VC
 - map the LCN from the X.25 VC over which the TL1 message is received to the destination NE's NSAP
 - determine the appropriate association to be used
 - forward the TL1 message over the appropriate association to the destination NE.

Page 8-27

- R8-88** [858] An NE shall send responses to TL1 commands using the same Application association to the GNE on which the request from the OS was received.

Page 8-28

- R8-89** [859] When a remote NE needs to send an autonomous TL1 message to a particular OS, the NE shall send the message to the GNE over the association established with the appropriate ACI (tl1Maintenance, tl1MemoryAdministration, or tl1Test).

Page 8-28

- R8-90** [860] To support remote NEs that need to send autonomous messages to a particular OS, the GNE shall support the following ACIs on the NE/NE interface in order to direct autonomous messages to the correct OS.

- ... tl1Maintenance OBJECT IDENTIFIER ::= {1 3 17 104 10 3 tl1Maintenance(2)}
- ... tl1MemoryAdministration OBJECT IDENTIFIER ::= {1 3 17 104 10 3 tl1MemoryAdministration(3)}
- ... tl1Test OBJECT IDENTIFIER ::= {1 3 17 104 10 3 tl1Test(4)}

Page 8-28

- R8-91** [861] The GNE shall be capable of establishing associations with subtending NEs for communication between a subtending NE and an OS using the following ACI:
- ... t11PeerComm OBJECT IDENTIFIER ::= {1 3 17 104 10 3
t11PeerComm(1)}
- Page 8-29
- R8-92** [862] The GNE shall maintain static information to map the X.25 address of each OS to its related ACI. This static information shall be provisionable.
- Page 8-29
- R8-93** [863] When an association is established between a GNE and a subtending NE for OS-NE communications, the GNE shall dynamically relate the association with the appropriate OS's X.25 virtual circuit.
- Page 8-29
- R8-94** [864] For NE-to-OS messages, the GNE shall route the TL1 message over an X.25 VC to the appropriate OS based on the association on which it was received.
- Page 8-29
- R8-95** [865] When an OS needs to communicate with an NE through a GNE, the GNE shall use or establish an association with the "target" NE using the appropriate ACI (t11Maintenance, t11MemoryAdministration, or t11Test). If the OS' X.121 address is not present as part of the GNE's mapping information, then the GNE shall use the t11PeerComm ACI.
- Page 8-29
- R8-96** [866] The OS shall initiate the X.25 VC to the GNE over the non-OSI TL1/X.25 interface.
- Page 8-29
- R8-97** [867] The GNE shall establish OSI Application associations with remote NEs.
- Page 8-30
- R8-98** [868] If the X.25 VC between the OS and the GNE goes down, either deliberately or through a communications failure, then the GNE shall take down the association(s) with the remote NE(s) using that X.25 VC.

- ... If multiple RNEs/VC are supported, then associations to several remote NEs will be taken down. If a single RNE/VC is supported, then an association to one remote NE is taken down.
Page 8-30
- R8-99** [869] All SONET NEs shall support the ES-IS protocol over the NE/NE operations communications interface as specified in Section C.5 of Appendix C.
Page 8-36
- O8-100** [870] All SONET NEs supporting the Redirect capability in the ES role of the ES-IS protocol should support the ISO 9542 Address Mask Generation function.
Page 8-36
- O8-101** [871] All SONET NEs supporting the Redirect capability in the IS role of the ES-IS protocol should be able to send the ISO 9542 PDU Address Mask field.
Page 8-36
- O8-102** [872] All SONET NEs supporting the ES or the IS role of the ES-IS protocol should be able to send and receive the ISO 9542 PDU Security field.
Page 8-36
- R8-103** [873] All SONET NEs supporting the IS-IS protocol shall do so in accordance with the protocol specifications in ISO 10589 and Appendix C.
Page 8-36
- O8-104** [874] SONET NEs supporting the IS-IS protocol should authenticate IS-IS PDUs based on passwords as specified in ISO 10589.
Page 8-36
- R8-105** [875] A SONET GNE shall maintain static information to map NSAP addresses of OSs to their corresponding X.121 addresses. This static information shall be provisionable locally or remotely.
Page 8-37
- R8-106** [876] All operations functions supported by a SONET NE via OS/NE communications shall be supported at the local craftsperson interface.
Page 8-38
-

-
- R8-107** [877] Any features available exclusively at the local craftsperson interface shall be clearly identified in supporting documentation.
Page 8-38
- R8-108** [878v2] The workstation shall provide the craftsperson a command-line mode of operation (or interaction) as defined in TR-TSY-000824.
Page 8-38
- R8-109** [879] The command-line interface between the craftsperson and the workstation shall conform to TL1 requirements in GR-831-CORE, *OTGR Section 12.1: Operations Application Messages – Language for Operations Application Messages*, which is based on User System Language (USL) of TR-TSY-000825, *OTGR Section 10.A: User System Interface – User System Language*.
Page 8-38
- R8-110** [880v2] The user interface requirements specified in GR-826-CORE, *OTGR Section 10.2: User Interface Generic Requirements for Supporting Network Element Operations*, shall be supported (e.g., support of a graphical user interface).
Page 8-38
- R8-111** [881v2] A WS/NE interface shall be provided in accordance with TR-TSY-000824.
Page 8-39
- O8-112** [1039] It is an objective that SONET NEs support the LAN-based interface defined by Section 4 of SIF-009-1997.
Page 8-39
- R8-113** [1040v2] SONET NEs shall support the WS/Remote NE interface defined by Section 5 of SIF-009-1997.
Page 8-39
- R8-114** [882] When a SONET NE supports TL1/OSI on the NE-NE interface (DCC or LAN), the NE shall also support TARP on the NE-NE interface according to the requirements of Sections 8.7 and C.8.
Page 8-39
- R8-115** [1041] For all TARP-related TID/NET mappings, case shall be ignored for the TIDs.
Page 8-39
-

-
- R8-116** [883v2] When an NE supports TARP and has multiple NETs associated with the same TID, the NE shall designate one NET as the NET to be used for mapping purposes for all TARP-related TID/NET mappings.
Page 8-40
- R8-117** [884] TARP PDUs shall be carried as ISO 8473 (CLNP) Data (DT) PDUs.
Page 8-40
- R8-118** [885] When TARP PDUs are sent, the following constraints shall apply to the header of the CLNP DT PDU:
- ... a. the PDU Lifetime field shall be set to a value of 25000 milliseconds
 - ... b. the Segmentation Permitted flag shall be set to a value of one (1) indicating that segmentation is permitted
 - ... c. the Error Report flag shall be set to a value of zero (0) indicating that discard of the PDU will not cause generation of an Error Report PDU.
- Page 8-40
- R8-119** [886] The TARP PDU fields shown in Table 8-1 shall be supported and sent in the order shown by Table 8-1 (starting with tar-lif).
Page 8-40
- R8-120** [1042] The URC bit shall be ignored upon receipt of TARP PDUs.
Page 8-42
- R8-121** [887] The NE shall process address resolution requests (from a higher layer application in the NE) to find the NET that matches a given TID, according to the procedure given in Section 8.7.4.1.
Page 8-43
- R8-122** [888] The NE shall process address resolution requests (from a higher layer application in the NE) to find the TID that matches a given NET, according to the procedure given in Section 8.7.4.2.
Page 8-43
- R8-123** [889] When a TID or Protocol Address change occurs at an NE, the NE shall notify other NEs of this change according to the procedure given in Section 8.7.4.3.
Page 8-43
-

-
- R8-124** [890] The NE shall provide the function of a TARP processor that is capable of originating and receiving TARP PDUs for all five TARP Types according to the descriptions given throughout Section 8.7.3. Page 8-45
- R8-125** [891v2] Each time an NE originates a TARP PDU, the NE shall increment the tar-seq field. The range of the tar-seq field shall be 0 to 65,535. If the tar-seq field reaches 65,535 (or if the NE is reset), a TARP Type 4 PDU shall be sent with a tar-seq field equal to zero, and the next TARP PDU shall be sent with the tar-seq field equal to 1. A zero value will notify all other NEs that a reset has occurred. Page 8-45
- CR8-126** [892] Whenever tar-seq is reset to zero, a TARP Type 4 PDU may be generated even if the TID and/or network address has not changed. Page 8-45
- R8-127** [1043] Inclusion of the tar-tor field in TARP Type 1 or Type 2 PDUs is optional (at the PDU sender's discretion). The receiver of TARP Type 1 or Type 2 PDUs shall ignore the contents of the tar-tor field (if present) and shall be capable of correctly processing the PDUs regardless of whether or not the tar-tor field is present. Page 8-46
- R8-128** [893] All NEs supporting an IS function shall maintain a circular (first-in first-out) TARP Loop Detection Buffer. Page 8-49
- R8-129** [1012] All NEs supporting an IS function shall maintain a LDB Entry Timer for each LDB entry for which tar-seq = zero. The timer shall be settable within a range of 1 to 10 minutes. The default value shall be 5 minutes. Page 8-49
- R8-130** [894] The LDB Flush Timer shall be settable within a range of 0 to 1440 minutes. The default value shall be 5 minutes. Page 8-49
- R8-131** [895] The NE shall allow TARP propagation to be selectively disabled by link/adjacency. Page 8-50
-

- R8-132** [896v2] The TARP PDU fields listed in Table 8-4 shall be provisionable in accordance with the default values and ranges specified by Table 8-4. The values of other TARP PDU fields shall not be provisionable.
Page 8-50
- R8-133** [897] The NE shall allow the value of all TARP timers (as shown in Table 8-3) to be provisionable.
Page 8-50
- R8-134** [898] The NE shall be capable of displaying (via the local WS at a minimum) the TDC, the LDB, and the TARP Sequence Number in use.
Page 8-50
- R8-135** [899] The NE shall provide a manual flush capability for the TDC and the LDB.
Page 8-51
- R8-136** [900v2] The NE shall allow manual provisioning of entries for the TDC and the LDB.
Page 8-51
- R8-137** [901] The NE shall allow the disabling of any of the following: all TARP functions, TARP propagation functions, TARP origination functions, or the TDC.
Page 8-51
- R8-138** [902] The NE shall allow TARP requests to be manually generated.
Page 8-51

Appendix B: Fiber Optic Transmission System Design Worksheets

This appendix contains Worksheets 1 through 4, which are forms that can be used in gathering fiber optic transmission system design information. Worksheet 1 consists of general terminal equipment information the supplier provides. Worksheets 2 and 3 consist of terminal equipment parameters under normal operating and short-term emergency conditions, respectively, that the supplier and system integrator provide. Worksheet 4 summarizes the fiber optic cable transmission parameters to be provided by the system integrator and by the suppliers.

Some criteria are distance-related. If a supplier provides a complete system (NE plus cable), then the average regenerator spacing may be used in verifying distance-related criteria. If a supplier provides only NE (and is not given any specific distance information by the BCC), then the supplier is to show that the criteria are satisfied as the number of intermediate regenerators varies. For example, for certain criteria that are worded "up to 250 miles," the supplier is to verify that the criteria are satisfied for a number of intermediate regenerators (e.g., 10), spanning the range of applications up to 250 miles (assuming that the BCC provides the proper cable).

WORKSHEET 1:

GENERAL INFORMATION

PAGE 1 OF 2

Supplier-Provided Information

System Information

Terminal Equipment Identification _____
Optical Line Rate (Mb/s) _____

Transmitter Information

General:

Identification _____
Optical Device Temperature Controller (e.g., TEC) _____
FDA Classification (e.g., Class I, Class II) _____
Product Change Designation (e.g., issue, revision) _____

Optical Source:

Type of Device (e.g., Laser, LED, etc.) _____
Material Composition of Source (e.g., In GaAs) _____
Generic Device Structure (e.g., DFB) _____

Transmitter Connector:

Manufacturer _____
Type (e.g., Biconic, FC) _____
Model Number _____
Classification (Multimode, Single-mode) _____
Mating Connector Model Number _____

Transmitter Pigtail:

General Fiber Type _____
Class of Fiber _____
Mode Field Diameter _____ μm

Receiver Information:

General:

Identification _____
Optical Device Temperature Controller (e.g., TEC) _____
Product Change Designation (e.g., issue, revision) _____

Optical Detector:

Type of Device (e.g., PIN, APD) _____
Material Composition of Detector (e.g., Ge, Si) _____

WORKSHEET 1:

GENERAL INFORMATION

PAGE 2 OF 2

Receiver Connector:

Manufacturer
Type (e.g., Biconic, FC)
Model Number
Classification (Multimode, Single-mode)
Mating Connector Model Number

Receiver Pigtail:

General Fiber Type
Class of Fiber
Mode Field Diameter

_____ μm

WDM Device Information:

Manufacturer
Model Number
Number of Channels

Attenuator Device Information:

Manufacturer
Model Number

Station Cable Information:

Manufacturer
Model Number
General Fiber Type
Class of Fiber

Interconnection Related Parameters

Mode Field Diameter:

Nominal: _____ μm ,
Tolerance: _____ μm

Cladding Diameter:

Nominal: _____ μm ,
Tolerance: _____ μm

Maximum Cladding Ovality:

_____ μm

Maximum Core/Cladding Concentricity Error:

_____ μm

Connector Information:

Connector Manufacturer
Connector Type (e.g., Biconic, FC)
Connector Model Number
Connector Classification (Multimode, Single-mode)

WORKSHEET 2: TERMINAL EQUIPMENT PARAMETERS PAGE 1 OF 2
STANDARD OPERATING CONDITIONS —
WORST-CASE VALUES

Supplier Provided Information

Environment

Room Ambient Temperature Range _____ °C (°F)
Relative Humidity Range _____ %

Transmitter:

Central Wavelength Measurement Method _____ Power Weighted _____ Peak
Central Wavelength Range: $\lambda_{Tmin} =$ _____
 $\lambda_{Tmax} =$ _____

Transmitter Power: $P_T =$ _____

Receiver:

BER	P_{R1} [dBm]	R_{max} [dBm]
10^{-11}	_____	_____
10^{-10}	_____	_____
10^{-9}	_____	_____
10^{-8}	_____	_____
10^{-7}	_____	_____
10^{-6}	_____	_____

Transceiver Specifications:

Maximum Dispersion	$DSR_{max1} =$ _____ ps/nm
Dispersion Power Penalty @ $BER^1 =$	$P_{D1} =$ _____ dB
Maximum Optical Reflection	$OR_{max} =$ _____ dB
Reflection Power Penalty @ $BER^{11} =$	$R_p =$ _____ dB

1. Parameters must be specified at 10^{-10} BER. A BCC may request the parameters at other BER values.

WORKSHEET 2: TERMINAL EQUIPMENT PARAMETERS PAGE 2 OF 2
STANDARD OPERATING CONDITIONS —
WORST-CASE VALUES

Attenuator Device:

Insertion Loss:

Attenuator Reflectance:

$U_{att} =$ _____ dB

$OR_{att} =$ _____ dB

WDM Device:

WDM Loss:

$U_{WDM} =$ _____ dB

Station Cable:

Loss:

Cutoff Wavelength:

$U_{SM} =$ _____ dB/km

$\lambda_{cc} =$ _____ nm

Connectors:

Loss:

Connector Variation:

Connector Reflectance:

$U_{con} =$ _____ dB

$\Delta U_{con} =$ _____ dB

$OR_{con} =$ _____ dB

System Integrator Provided Information

Nominal Central Wavelength:

Safety Margin:

$\lambda_{Tnom} =$ _____

$M =$ _____ dB

Station Cable Length:

Transmitter

_____ km

Receiver

_____ km

Total

$l_{SM} =$ _____ km

Number of Connectors:

Transmitter²

Receiver³

Total

$N_{con} =$ _____

2. Excluding transmitter module connector.

3. Excluding receiver module connector.

WORKSHEET 3: TERMINAL EQUIPMENT PARAMETERS PAGE 1 OF 2
EXTENDED OPERATING CONDITIONS —
WORST-CASE VALUES

Supplier Provided Information

Environment

Room Ambient Temperature Range _____ °C (°F)
Relative Humidity Range _____ %

Transmitter:

Central Wavelength Measurement Method _____ Power Weighted _____ Peak
Central Wavelength Range: $\lambda_{Tmin} =$ _____
 $\lambda_{Tmax} =$ _____

Transmitter Power: $P_T =$ _____

Receiver:

BER	P_{R2} [dBm]	R_{max} [dBm]
10 ⁻¹¹	_____	_____
10 ⁻¹⁰	_____	_____
10 ⁻⁹	_____	_____
10 ⁻⁸	_____	_____
10 ⁻⁷	_____	_____
10 ⁻⁶	_____	_____

Transceiver Specifications:

Maximum Dispersion $D_{SRmax2} =$ _____ ps/nm
Dispersion Power Penalty $P_{D2} =$ _____ dB
@ BER⁴ =
Maximum Optical Reflection $OR_{max} =$ _____ dB
Reflection Power Penalty $R_p =$ _____ dB
@ BER =

4. Parameters must be specified at 10⁻¹⁰ BER. A BCC may request the parameters at other BER values.

WORKSHEET 3: TERMINAL EQUIPMENT PARAMETERS PAGE 2 OF 2
EXTENDED OPERATING CONDITIONS —
WORST-CASE VALUES

Attenuator Device:

Insertion Loss:

$U_{att} =$ _____ dB

Attenuator Reflectance:

$OR_{att} =$ _____ dB

WDM Device:

WDM Loss:

$U_{WDM} =$ _____ dB

Station Cable:

Loss:

$U_{SM} =$ _____ dB/km

Cutoff Wavelength:

$\lambda_{cc} =$ _____ nm

Connectors:

Loss:

$U_{con} =$ _____ dB

Connector Variation:

$\Delta U_{con} =$ _____ dB

Connector Reflectance:

$OR_{con} =$ _____ dB

System Integrator Provided Information

Nominal Central Wavelength:

$\lambda_{Tnom} =$ _____

Safety Margin:

$M =$ _____ dB

Station Cable Length:

Transmitter

Receiver

Total

_____ km

_____ km

$l_{SM} =$ _____ km

Number of Connectors:

Transmitter⁵

Receiver⁶

Total

$N_{con} =$ _____

5. Excluding transmitter module connector.

6. Excluding receiver module connector.

**WORKSHEET 4: CABLE TRANSMISSION PARAMETERS PAGE 1 OF 1
FOR A SPECIFIC APPLICATION**

System Integrator Provided Information

Application: Underground____ Aerial____ Buried____
Temperature Range: _____°C to _____°C
Cabled Fiber Reel Length: l_R = _____km
Nominal Central Wavelength: λ_{Tnom} = _____nm
Central Wavelength Range: λ_{Tmin} = _____nm to λ_{Tmax} = _____nm

Splices:

Type of splice: _____
Splice loss at 23°C: U_S = _____dB/splice
Maximum additional splice loss due to temperature variation: U_{ST} = _____dB/splice

Supplier Provided Information

Cable Designation: _____
Maximum Cable Cutoff Wavelength (per EIA/TIA-455-170 with cable deployment conditions shown in Figure 4-8): λ_{cc} = _____nm
Cable Loss at λ_{Tnom} and 23°C: U_c = _____dB/km
maximum additional loss at 23°C due to wavelength variation: U_λ = _____dB/km
maximum additional loss at λ_T due to temperature variation: U_{CT} = _____dB/km

Dispersion Parameters:

Zero-Dispersion Wavelength Range: λ_{omin} = _____nm to λ_{omax} = _____nm
Maximum Zero-Dispersion Slope: S_{0max} = _____ps/(nm²·km)
Worst Case Chromatic Dispersion Over Wavelength Range: D_{max} = _____ps/(nm·km)

Splice Related Parameters:

Mode Field Diameter: Nominal: _____μm, Tolerance: _____%
Cladding Diameter: Nominal: _____μm, Tolerance: _____%
Maximum Cladding Ovality: _____%
Maximum Core/Cladding Concentricity Error: _____μm

Appendix C: SONET Operations Communications Lower Layers Protocol Profile

C.1 Introduction

The SONET profile specifies the implementation requirements for the SONET Operations Communications protocol stack. This lower-layer profile provides a set of tables for each of the protocols used in Layers 2 through 4 (i.e., LAPD (for the DCC), LLC (for the LAN), CLNP, IS-IS, ES-IS, and TP4). The structure and content of the tables within this Appendix are the same as those for each of the PICS proforma contained in the corresponding ISO/CCITT Standard. An additional Bellcore profile column has been added that specifies the Bellcore requirements for a given item. The support column in the tables can be used:

- by the protocol implementor, as a checklist to ensure more complete compliance with the Bellcore profile
- by the supplier and acquirer, or potential acquirer, of the implementation, as a detailed indication of the capabilities of the implementation
- by the user - or potential user - of the implementation, as a basis for initially checking the possibility of interworking with another implementation (note that, while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible profiles)
- by a protocol analyzer as the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation.

C.1.1 Source Documents

When defining the SONET Lower Layers Profile, the following source documents were used:

C.1.1.1 Base Standards

CCITT Recommendation Q.921:1988: *ISDN User-Network Interface - Data Link Layer Specification*

ISO/IEC 8073:1988: *Information processing systems - Open Systems Interconnection - Connection oriented transport protocol specification*

ISO/IEC 8073 AD2:1989: *Information processing systems - Open Systems Interconnection - Connection oriented transport protocol specification - Addendum 2: Class four operation over connectionless network service*

ISO/IEC 8473:1988: *Information processing systems - Data communications - Protocol for providing the connectionless-mode network layer service*

ISO/IEC 8802-2:1989: *Information processing systems - Open Systems Interconnection - Local area networks - Part 2: Logical link control*

ISO/IEC 8802-3:1990: *Information processing systems - Open Systems Interconnection - Local area networks - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*

ISO/IEC 9542:1988: *Information processing systems - Telecommunications and information exchange between systems - End system to Intermediate system routing exchange protocol for use in conjunction with protocol for providing the connectionless-mode Network Service (ISO 8473)*

ISO/IEC 10589:1992: *Information technology - Telecommunications and information exchange between systems - Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

C.1.1.2 PICS Proforma

CCITT Recommendation Q.921:1988: *PICS Proforma for ISDN D-Channel, Layer 2 Basic Rate Access, User Side*

ISO/IEC 8073:1992: *Information technology - Telecommunications and information exchange between systems - Open Systems Interconnection - Protocol for providing the connection-mode transport service - Annex C*

ISO/IEC DIS 8473-1:1993: *Information technology - Protocol for providing the connectionless-mode network service - Annex A*

ISO/IEC 8802-2:1989/Amd 3 - *Information Processing Systems - Local Area Networks - Part 2: Logical Link Control - Amendment 3: Conformance Requirements*

ISO/IEC 9542:1988: *Information processing systems - Telecommunications and information exchange between systems - End system to Intermediate system routing exchange protocol for use in conjunction with protocol for providing the connectionless-mode network service (ISO 8473) - Annex A*

ISO/IEC 10589:1992: *Information technology - Telecommunications and information exchange between systems - Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) - Annex A*

C.1.1.3 International Standardized Profiles

ISO/IEC ISP 10608-1:1992: *Information technology - International Standardized Profile
TAnnnn - Connection-mode Transport Service over Connectionless-mode Network
Service - Part 1: General overview and subnetwork-independent requirements*

ISO/IEC ISP 10608-2:1992: *Information technology - International Standardized Profile
TAnnnn - Connection-mode Transport Service over Connectionless-mode Network
Service - Part 2: TA51 profile including subnetwork-dependent requirements for
CSMA/CD Local Area Networks (LANs)*

C.1.1.4 American National Standards for Telecommunications

ANSI T1.204-1993, *Operations, Administration, Maintenance and Provisioning
(OAM&P) - Lower Layer Protocols for Telecommunications Management Network
(TMN) Interfaces between Operations Systems and Network Elements*

C.1.1.5 Bellcore Requirements

Section 8 of this GR
GR-828-CORE

C.1.2 Notations Used in the SONET Lower Layer Profile

C.1.2.1 Status Symbols

Each PICS proforma has a set of symbols used in the status column of the PICS. These symbols appear (or are referred to) at the beginning of each protocol profile. |

C.1.2.2 Profile Symbols

The profile symbols are the set of symbols used in the Bellcore profile column. These symbols present Bellcore's view regarding a given item in the PICS proforma.

M mandatory support is required

M:n mandatory support is required given the condition n

M(N) mandatory support is required as clarified by additional comment N

i the item is out of scope for this profile

i(N)	the item is out of scope for this profile as clarified by additional comment N
X	the item is excluded (i.e., prohibited) for this profile
X(N)	the item is excluded (i.e., prohibited) for this profile as clarified by additional comment N
n/a	the item is not applicable for this profile
O.n	the item is optional, but, if chosen, support is required for either at least one or only one of the options in the group labeled by the same numeral <n>
ENE	the status following this symbol applies only when the NE is providing End NE functionality as described in Section 8.1.4.
INE	the status following this symbol applies only when the NE is providing Intermediate NE functionality as described in Section 8.1.3.
GNE	the status following this symbol applies only when the NE is providing Gateway NE functionality as described in Section 8.1.2.

When an item is designated as 'i' (out of scope), implementation of that item is not precluded; however, for the purposes of the SONET Lower Layers Profile, the implementation or non-implementation of that item is ignored.

When an item is designated as 'n/a' (not applicable), that item is not used for SONET Operations Communications Applications. For the purpose of the SONET Lower Layers Profile the item is not used.

There are instances in the Bellcore profile where Bellcore has declared an item to be prohibited, while it is mandatory in the base standard. Bellcore is not stating that a supplier's OSI implementation should not have this feature; rather Bellcore is stating that the given item's use in SONET applications is prohibited.

Similarly, there are instances in the Bellcore profile where the comment "see note #", where # is a given note number, appears. These instances signify Bellcore objectives, and should be considered as such.

All conditionally mandatory functions in the protocol profile that are **always true** have been modified. The profile column no longer includes the conditional, and is made strictly mandatory. This is to reduce the confusion associated with searching to determine whether the condition is true. All truly conditional functions (where the conditional is out of scope for the profile) are left conditional in the profile column.

The SONET protocol profiles follow the standardized PICS proformas as closely as possible. There are places where there have been editorial modifications to the Standard PICS, allowing the profile to be more clearly understood. Some sections of the standardized PICS proformas have been deleted from the SONET protocol profiles. The sections of the standardized PICS proformas that are not contained within this document are to be viewed as either out of scope of the SONET protocol profile or not applicable for SONET applications. For example, the SONET DCC uses only a subset of the LAPD capabilities.

Thus much of the unused LAPD capabilities (e.g., Terminal End-point Identifier [TEI] management) are not applicable to SONET applications and are not included in this profile. Therefore, there may be parameters in each protocol profile that are out-of-sequence.

C.1.2.3 Support Symbols

The support symbols are the set of symbols used in the Support column. These symbols present the supplier's or implementor's view of the Bellcore profiles.

- Y** yes, the feature has been implemented;
- N** no, the feature has not been implemented;
- n/a** not applicable.

C.1.2.4 References

The reference columns in the following tables contain the reference or references to the material that specifies the item in the main body of corresponding base ISO/IEC and CCITT (Blue Book) international standards.

C.2 LAPD SONET Protocol Profile

The following SONET profile specifies the implementation requirements for the Link Access Procedure on the D-Channel (LAPD). The protocol profile is based on the PICS proforma contained in Annex E of CCITT Recommendation Q.921 for Basic Rate, with an additional Bellcore profile column.

C.2.1 Notations

C.2.1.1 Abbreviations

APPX	Appendix
DCC	Data Communications Channel
DLCI	Data Link Connection Identifier, DLCI=(SAPI,TEI)
DLE	Data Link Entity
FR	prefix for Index of the Frames group
IUT	Implementation Under Test
NE	Network Element
PC	prefix for the Index number of the Protocol Capabilities group
SAPI	Service Access Point Identifier
SP	prefix for the Index number of the System Parameter group
TEI	Terminal End-point Identifier

C.2.1.2 Status Symbols

M	mandatory
O	optional
O.<n>	optional, but, if chosen, support is required for either at least one or only one of the options in the group labeled by the same numeral <n> is required
X	prohibited

C.2.1.3 Additional Symbols

<r>	receive aspects of an item
<s>	send aspects of an item

C.2.2 Protocol Capabilities (PC)

Index	Protocol Feature	Status	Reference	Profile	Support
PC1.1	Is the <i>NE</i> of the non-automatic TEI assignment category?	O.1	3.3.4.2	M	
PC1.2	Is the <i>NE</i> of the automatic TEI assignment category?	O.1	3.3.4	i	
PC2	Does the <i>NE</i> support the broadcast data link?	M	5.2	M:1	
PC4	Does the <i>NE</i> support data link monitor function?	O	5.10	M	
PC5	Does the <i>NE</i> support reject retransmission procedure?	O	3.6.7, 5.8.1 Appx. I	i	
PC6.1	Does the DLE support automatic negotiation of data link layer parameters?	O.2	Appx. IV	i	
PC6.2	Does the DLE support internal parameter initialization?	O.2	5.4	M	
PC7	Does the <i>NE</i> permit concurrent LAPB data link connection within the D-channel?	O	2.3	n/a	
Service Access Point Identifier (SAPI)					
PC8	If the <i>NE</i> supports Layer 3 call control procedures, is SAPI=0 supported?	M	3.3.3	M:2	
PC9	If the <i>NE</i> supports X.25 Layer 3 packet procedures on the DCC, is SAPI=16 supported?	M	3.3.3	M:2	
PC10	Is SAPI=63 supported?	M	3.3.3	M:2	

O.1 = Support of at least one of these items is required.

O.2 = Support of at least one of these items is required.

Notes:

- 1 Support of the Broadcast Data Link functionality is mandatory, but the functionality is not used for SONET DCC applications.
- 2 SAPI subfields of the LAPD address field for LAPD D-channel applications do not apply to the SONET DCC applications. SAPI value of 62 is used for SONET DCC applications. The need to reserve additional SAPI values for specific purposes is for further study.

Index	Protocol Feature	Status	Reference	Profile	Support
PC11.1	Does the implementation support the association of a given TEI with all SAPs which the <i>NE</i> supports?	O	3.3.4,5.3.1 (Q.920 3.4.3)	i	
PC11.2	If the <i>NE</i> is an X.31 type of packet mode terminal equipment, is a given TEI for point-to-point data link connection (<127) associated with all SAPs which the <i>NE</i> supports?	M	3.3.4,5.3.1 (Q.920 3.4.3)	M:3	
PC12	Does the implementation support modulus 128 for frames numbering?	M	3.5.2.1, 5.5.1	M	
Peer-to-Peer Procedures					
<i>Unacknowledged Information Transfer</i>					
PC13	Does the <i>NE</i> support UI-command?	M	5.2.2	M	
PC14	Is the P/F bit set to 0?	M	5.1.1	M	
TEI Management					
PC15	Does the <i>NE</i> transmit management entity messages in UI frames with DLCI=(63,127)	O.3	5.3.1	i(4)	
Establishment and Release of Multiple Frame Operation					
PC32	Does the <i>NE</i> support multiple frame operation?	M	5.5	M	
<i>Does the DLE initiate multi-frame establishment</i>					
PC33.1	a) immediately after TEI assignment?	O.7	5.5	M	
PC33.2	b) when there is an incoming or an outgoing call?	O.7	5.5	n/a	
PC34.1	c) Does the DLE remain in TEI Assigned state when the multiple frame operation is released?	O.8	5.5.3	M	
PC34.2	d) Does the DLE initiate immediate re-establishment when the multiple frame operation is released?	O.8	5.5.3	M	

O.3 = Support of at least one of these items is required.

O.7 = Support of at least one of these items is required.

O.8 = Support of at least one of these items is required.

Notes:

3 See note 2.

4 Support of the TEI Management functionality is conditionally optional, but the functionality is not used for SONET DCC applications.

C.2.3 Frames - Protocol Data Units (FR)

Index	Protocol Feature	Status	Reference	Profile	Support
Frame Format					
FR1	Format A	M	2.1	M	
FR2	Format B	M	2.1	M	
Flag Sequence					
FR3	Opening flag	M	2.2	M	
FR4	Closing flag	M	2.2	M	
Address Field					
FR5	Two octets	M	2.3	M	
FR6	If the DLE permits concurrent LAPB data link connection with the D-channel, is the one octet address field recognized?	M	2.3	n/a	
Control Field					
<i>Unacknowledged operation</i>					
FR7	Single octet	M	2.4	M	
<i>Multiple frame operation</i>					
FR8	Two octets	M	2.4	M	
FR9	Single octet (unnumbered frame)	M	2.4	M	
Order of Bit Transmission					
FR10	Ascending numerical order	M	2.8.2	M	
Field Mapping Convention					
FR11	Lowest bit number = Lowest order value	M	2.8.3	M	
<i>Do all transmitted frames contain the following fields?</i>					
FR12.1	- Flag	M	2.2	M	
FR12.2	- Address	M	2.3	M	
FR12.3	- Control	M	2.4	M	
FR12.4	- FCS	M	2.7	M	
FR13	Is the NE capable of accepting the closing flag as the opening flag of the next frame?	M	2.2	M	
FR14	Does the NE generate a single flag as above?	O	2.2	i	
FR15	Does the NE ignore one flag, or two or more consecutive flags that do not delimit frames?	M	2.2	M	
FR16	Are all invalid frames discarded and no action taken?	M	2.9	M	
FR17	Are seven or more contiguous I bits interpreted as an abort and the associated frames ignored?	M	2.10	M	
FR18	If the NE supports the automatic negotiation of data link layer parameters, does the NE support XID frames?	M	Appx IV	n/a	

C.2.4 System Parameters (SP)

Index	Protocol Feature	Status	Reference	Profile	Supported Range/ Default	Support
	<i>If the DLE supports multiple frame operation</i>					
SP1	Retransmission time (T200)	M	5.9.1	M	0.2 to 20 seconds/ 0.2 seconds	
SP2	Maximum number of retransmission	M	5.9.2	M	2 to 16 frames/ 3 frames	
	<i>Maximum number of octets in information field</i>					
SP3	for SAP supporting signaling	M	5.9.3	n/a	n/a	
SP4	for SAP supporting packet on the DCC	M	5.9.3	M	512 or greater octets/ 512 octets, see note 5	
	<i>Maximum number of outstanding I frames (k)</i>					
SP5	for SAP supporting basic access signaling	M	5.9.5	n/a	n/a	
SP6	for SAP supporting basic access packet on the DCC	M	5.9.5	M	1 to 127 frames/ 7 frames	
	<i>If the NE supports the data link monitor function,</i>					
SP9	Maximum time allowed without frames being exchanged (T203)	M	5.9.8	M	4 to 120 seconds/ 10 seconds	

Notes:

- 5 For both I-Frame and UI-Frame operation

C.3 LLC Protocol Profile

The following SONET profile specifies the implementation requirements for the Logical Link Control Protocol (ISO 8802-2:1989). The protocol profile is based on the PICS proforma contained in Annex C of ISO 8802-2:1989/Amd. 3, with an additional Bellcore profile column.

C.3.1 Abbreviations and Special Symbols

C.3.1.1 Status Symbols

M	mandatory
O	optional
O.<n>	optional, but support of at least one of the group of options labeled by the same numeral <n> is required
X	prohibited
<item>	conditional symbol, status is dependent on the support marked for <item>

C.3.1.2 Item References

The following is a list of item references used in the PICS proforma:

CLS	Class of LLC supported
UI	UI PDUs
XID	XID PDUs
TES	TEST PDUs
UIT	UI PDUs transmitted
XDT	XID PDUs transmitted
TST	TEST PDUs transmitted
UIR	UI PDUs received
XDR	XID PDUs received
TSR	TEST PDUs received
MIS	Miscellaneous

C.3.2 Claimed Conformance to ISO 8802-2:1989/Amd.1, Amd.2 and Amd.4

Item	Protocol Feature	References	Status	Profile	Support
CLS1a	Is Class 1 LLC supported?	4.2	O.1	M	
CLS1b	Are LLC Type 1 procedures supported?	4.2	CLS1a:M	M	
CLS2a	Is Class 2 LLC supported?	4.2	O.1	i	
CLS2b	Are LLC Type 1 and Type 2 procedures supported?	4.2	CLS2a:M	i	
CLS3a	Is Class 3 LLC supported?	4.2	O.1	i	
CLS3b	Are LLC Type 1 and Type 3 procedures supported?	4.2	CLS3a:M	i	
CLS4a	Is Class 4 LLC supported?	4.2	O.1	i	
CLS4b	Are LLC Type 1, Type 2 and Type 3 procedures supported?	4.2	CLS4a:M	i	

C.3.3 LLC Type 1 Operation - Unacknowledged Connectionless Mode

C.3.3.1 LLC Type 1 - Supported PDU Types

Item	Protocol Feature Supported PDU types	References	Status	Profile	Support
UI/1	UI_CMD supported on transmission	6.1, 6.5.1	M	M	
UI/2	UI_CMD supported on receipt	6.1, 6.5.2	M	M	
XID/3	XID_CMD supported on transmission	6.6	O	i	
XID/4	XID_CMD supported on receipt	6.6	M	M	
XID/5	XID_RSP supported on transmission	6.6	M	M	
XID/6	XID_RSP supported on receipt	6.6	XID/3:M	XID/3:M	
TES/7	TEST_CMD supported on transmission	6.7	O	i	
TES/8	TEST_CMD supported on receipt	6.7	M	M	
TES/9	TEST_RSP supported on transmission	6.7	M	M	
TES/10	TEST_RSP supported on receipt	6.7	TES/7:M	TES/7:M	

C.3.3.2 LLC Type 1 - Supported Parameters in PDUs on Transmission

Item	Protocol Feature Supported PDU types	References	Status	Profile	Support
UIT/11	UI_CMD - DSAP address	6.2	M	M	
UIT/12	UI_CMD - SSAP address	6.2	M	M	
UIT/13	UI_CMD - P-bit = 0	6.3	M	M	
UIT/14	UI_CMD - information	3.3	O	M	
XDT/15	XID_CMD - DSAP address	6.2, 6.6	XID/3:M	XID/3:M	
XDT/16	XID_CMD - SSAP address	6.2, 6.6	XID/3:M	XID/3:M	
XDT/17	XID_CMD - P-bit = 1	6.3	XID/3:O.2	XID/3:O.2	
XDT/18	XID_CMD - P-bit = 0	6.3	XID/3:O.2	XID/3:O.2	
XDT/19	XID_CMD - information	5.4.1.1.2, 6.6	XID/3:M	XID/3:M	
XDT/20	XID_RSP - DSAP address	6.2, 6.6	M	M	
XDT/21	XID_RSP - SSAP address	6.2, 6.6	M	M	
XDT/22	XID_RSP - F-bit = P-bit	6.3	M	M	
XDT/23	XID_RSP - information	5.4.1.2.1, 6.6	M	M	
TST/24	TEST_CMD - DSAP address	6.2	TES/7:M	TES/7:M	
TST/25	TEST_CMD - SSAP address	6.2	TES/7:M	TES/7:M	
TST/26	TEST_CMD - P-bit = 1	6.3	TES/7:O.3	TES/7:O.3	
TST/27	TEST_CMD - P-bit = 0	6.3	TES/7:O.3	TES/7:O.3	
TST/28	TEST_CMD - information	5.4.1.1.3, 6.7	TES/7:O	i	
TST/29	TEST_RSP - DSAP address	6.2	M	M	
TST/30	TEST_RSP - SSAP address	6.2	M	M	
TST/31	TEST_RSP - F-bit = P-bit	6.3	M	M	
TST/32	TEST_RSP - information	5.4.1.2.2, 6.7	M	M	

C.3.3.3 LLC Type 1 - Supported Parameters in PDUs on Receipt

Item	Protocol Feature Supported PDU types	References	Status	Profile	Support
UIR/33	UI_CMD - DSAP address	6.2	M	M	
UIR/34	UI_CMD - SSAP address	6.2	M	M	
UIR/35	UI_CMD - P-bit = 0	6.3	M	M	
UIR/36	UI_CMD - information	3.3	O	M	
XDR/37	XID_CMD - DSAP address	6.2, 6.6	M	M	
XDR/38	XID_CMD - SSAP address	6.2, 6.6	M	M	
XDR/39	XID_CMD - P-bit = 1	6.3	M	M	
XDR/40	XID_CMD - P-bit = 0	6.3	M	M	
XDR/41	XID_CMD - information	5.4.1.1.2, 6.6	M	M	
XDR/42	XID_RSP - DSAP address	6.2, 6.6	M	M:1	
XDR/43	XID_RSP - SSAP address	6.2, 6.6	M	M:1	
XDR/44	XID_RSP - F-bit = P-bit	6.3	M	M:1	
XDR/45	XID_RSP - information	5.4.1.2.1, 6.6	M	M:1	
TSR/46	TEST_CMD - DSAP address	6.2	M	M	
TSR/47	TEST_CMD - SSAP address	6.2	M	M	
TSR/48	TEST_CMD - P-bit = 1	6.3	M	M	
TSR/49	TEST_CMD - P-bit = 0	6.3	M	M	
TSR/50	TEST_CMD - information	5.4.1.1.3, 6.7	M	M	
TSR/51	TEST_RSP - DSAP address	6.2	M	M:2	
TSR/52	TEST_RSP - SSAP address	6.2	M	M:2	
TSR/53	TEST_RSP - F-bit = P-bit	6.3	M	M:2	
TSR/54	TEST_RSP - information	5.4.1.2.2, 6.7	TST/28:M	TST/28:M	

Notes:

- 1 Support is mandatory if XID_RSP on reception is supported (see XID/6).
- 2 Support is mandatory if TES_RSP on reception is supported (see TES/10).

C.3.3.4 LLC Type 1 - Miscellaneous

Item	Protocol Feature Supported PDU types	References	Status	Profile	Support
MIS/55	Do all transmitted PDUs contain an integral number of octets	3.3	M	M	
	If the following PDUs are received from the MAC sub-layer are they treated as invalid and ignored:				
MIS/56	-contains a non-integral number of octets	3.3.5	M	M	
MIS/57	-has a length less than 3 octets	3.3.5	M	M	

LLC Type 1 - Miscellaneous (Continued)

Item	Protocol Feature Supported PDU types	References	Status	Profile	Support
	Which of the following addresses are supported in the DSAP address field of UI PDUs				
MIS/58	-individual address	5.4.1.1.1	O.4	M	
MIS/59	-group address	5.4.1.1.1	O.4	M	
MIS/60	-global address	5.4.1.1.1	O.4	M	
MIS/61	-null address	5.4.1.1.1	O.4	M	
MIS/62	Is the address in the SSAP address field of a UI PDU the originator's individual address	5.4.1.1.1	M	M	
MIS/63	Are all UI PDU's transmitted as UI_CMD PDU's	6.5.1	M	M	
MIS/64	Are all UI_CMD PDU's transmitted with the P-bit = 0	6.5.1	M	M	
MIS/65	If a UI_CMD PDU is received with the P-bit = 1 is the PDU discarded?	6.3	O	i	
MIS/66	If a UI_RSP PDU is received is the frame discarded	6.5.2	M	M	
	Which of the following addresses are supported in the DSAP address field of XID_RSP PDUs				
MIS/73	-individual address	5.4.1.2.1	M	M	
MIS/74	-null address	5.4.1.2.1	M	M	
	Which of the following addresses are supported in the SSAP address field of XID_RSP PDUs				
MIS/75	-individual address	5.4.1.2.1	M	M	
MIS/76	-null address	5.4.1.2.1	M	M	
	Which of the following addresses are supported in the DSAP address field of TEST_RSP PDUs				
MIS/83	-individual address	5.4.1.2.2	M	M	
MIS/84	-null address	5.4.1.2.2	M	M	
	Which of the following addresses are supported in the SSAP address field of TEST_RSP PDUs				
MIS/85	-individual address	5.4.1.2.2	M	M	
MIS/86	-null address	5.4.1.2.2	M	M	
MIS/87	Is Duplicate Address Checking supported	6.9.2	O	i	
MIS/88	Is the ACK_TIMER function supported	6.9.2	MIS/87:M	MIS/87:M	
MIS/89	ACK_TIMER range			i	
MIS/90	Is the RETRY_COUNTER function supported	6.9.2	MIS/87:M	MIS/87:M	
MIS/91	RETRY_COUNTER range			i	
MIS/92	Is the XID_R_COUNTER function supported	6.9.2	MIS/87:M	MIS/87:M	

C.4 ISO 8473 Protocol Profile

The following SONET profile specifies the implementation requirements for the Connectionless-mode Network Layer Protocol (CLNP). The protocol profile is based on the PICS proforma contained in Annex A of ISO/IEC DIS 8473-1:1993, with an additional Bellcore profile column.

C.4.1 Notations

C.4.1.1 Status Symbols

M	mandatory
O	optional
O.<n>	optional, but support of at least one of the group of options labeled by the same numeral <n> is required
X	prohibited
<pred>	conditional-item symbol, including predicate identification
^	logical negation, applied to a conditional item's predicate
*	each item whose reference is used in a predicate or a predicate definition is indicated by an asterisk in the Item column

C.4.1.2 Additional Symbols

<r>	receive aspects of an item
<s>	send aspects of an item

C.4.2 Major Capabilities

Item	Capability	Reference	Status	Profile	Support
*ES	End system		O.1	M	
*IS	Intermediate system		O.1	M:1	
FL-r	<r> Full protocol	6	M	M	
FL-s	<s> Full protocol	6	M	M	
NSS-r	<r> Non-segmenting subset	5.2	M	M	
*NSS-s	<s> Non-segmenting subset	5.2	IS:M, ^IS:O	X(2)	
*IAS-r	<r> Inactive subset	5.2	ES:O	M:3	
*IAS-s	<s> Inactive subset	5.2	IAS-r:M, ^IAS-r:X	X(4)	

Notes: 1 Support is mandatory for IS NEs.

2 Implementation of this function is conditionally mandatory (for IS implementations) or conditionally optional (for ^IS implementations) in the base PICS, but its use is **prohibited** in this profile.

3 Received NPDUs encoded with the inactive subset are discarded.

4 Implementation of this function is conditionally mandatory (for IAS-r implementations) or conditionally optional (for ^IAS-r implementations) in the base PICS, but its use is **prohibited** in this profile.

C.4.3 End Systems

C.4.3.1 Applicability

The profile items in Section C.4.3 are applicable only to end system implementations; i.e., those in which item ES in Section C.4.2 is supported.

C.4.3.2 Supported Functions

Item	Function	Reference	Status	Profile	Support
ePDUC	PDU composition	6.1	M	M	
ePDUD	PDU decomposition	6.2	M	M	
ePHFA	Header format analysis	6.3	M	M	
ePDUL-s	<s> PDU lifetime control	6.4	M	M	
ePDUL-r	<r> PDU lifetime control	6.4	O	M	
eRout	Route PDU	6.5	M	M	
eForw	Forward PDU	6.6	M	M	
eSegm	Segment PDU	6.7	M	M	
eReas	Reassemble PDU	6.8	M	M	
eDisc	Discard PDU	6.9	M	M	
eErep	Error reporting	6.10	M	M	
eEdec-s	<s> Header error detection	6.11	M	M	
eEdec-r	<r> Header error detection	6.11	M	M	
*eSecu-s	<s> Security	6.13	O	i	
*eSecu-r	<r> Security	6.13	O	i	
*eCRR-s	<s> Complete route recording	6.15	O	i	
*eCRR-r	<r> Complete route recording	6.15	O	i	
*ePRR-s	<s> Partial route recording	6.15	O	i	
*ePRR-r	<r> Partial route recording	6.15	O	i	
*eCSR	Complete source routing	6.14	O	i	
*ePSR	Partial source routing	6.14	O	X	
*ePri-s	<s> Priority	6.17	O	i	
*ePri-r	<r> Priority	6.17	O	i	
*eQOSM-s	<s> QoS maintenance	6.16	O	M	
*eQOSM-r	<r> QoS maintenance	6.16	O	M	
*eCong-s	<s> Congestion notification	6.18	eQOSM-s:M	eQOSM-s:M, see note 5	
*eCong-r	<r> Congestion notification	6.18	O	see note 6	
*ePadd-s	<s> Padding	6.12	O	i	
ePadd-r	<r> Padding	6.12	M	M	
eEreq	Echo request	6.19	O	i(7)	
eEresp	Echo response	6.20	O	i(7)	
eSegS	Create segments smaller than necessary	6.8	O	i	

Notes:

- 5 Implementation of this function is conditionally mandatory (for eQOSM-s implementations) in the base PICS, but its use is a **Belcore objective**.
- 6 This is a **Belcore objective**. It is defined to be **out of scope** for this profile.
- 7 This is a new function introduced in the 1993 version of ISO/IEC 8473. To be compatible with ISO/IEC ISP 10608-1 (which is based on the 1988 version of ISO/IEC 8473), this new functionality is defined to be **out of scope** for this profile

C.4.3.3 Supported PDUs

Item	NPDU	Reference	Status	Profile	Support
eDT-t	DT (full protocol) transmit	7.7	M	M	
eDT-r	DT (full protocol) receive	7.7	M	M	
eDTNS-t	DT (non-segmenting) transmit	7.7	NSS-s:M	X	
eDTNS-r	DT (non-segmenting) receive	7.7	M	M	
eER-t	ER transmit	7.9	M	M	
eER-r	ER receive	7.9	M	M	
eIN-t	Inactive PDU transmit	7.8	IAS-s:M	X	
eIN-r	Inactive PDU receive	7.8	IAS-r:M	M(8)	
eERQ-t	ERQ transmit	7.10	eEreq:M	i(9)	
eERQ-r	ERQ receive	7.10	M	i(9)	
eERP-t	ERP transmit	7.11	eErsp:M	i(9)	
eERP-r	ERP receive	7.11	M	i(9)	

Notes: 8 See note 3.

9 See note 7. Note that implementation of this function may cause a backwards incompatibility with 1988 version of ISO/IEC 8473.

C.4.3.4 Supported Parameters

C.4.3.4.1 DT Parameters

Item	Parameter	Reference	Status	Profile	Support
edFxFt-s	<s> Fixed part	7.2	M	M	
edFxFt-r	<r> Fixed part	7.2	M	M	
edAddr-s	<s> Addresses	7.3	M	M	
edAddr-r	<r> Addresses	7.3	M	M	
edSeg-s	<s> Segmentation part	7.4	M	M	
edSeg-r	<r> Segmentation part	7.4	M	M	
edPadd-s	<s> Padding	7.5.2	ePadd-s:M	ePadd-s:M	
edPadd-r	<r> Padding	7.5.2	M	M	
edSecu-s	<s> Security	7.5.3	eSecu-s:M	eSecu-s:M	
edSecu-r	<r> Security	7.5.3	eSecu-r:M	eSecu-r:M	
edCRR-s	<s> Complete route recording	7.5.5	eCRR-s:M	eCRR-s:M	
edCRR-r	<r> Complete route recording	7.5.5	eCRR-r:M	eCRR-r:M	
edPRR-s	<s> Partial route recording	7.5.5	ePRR-s:M	ePRR-s:M	
edPRR-r	<r> Partial route recording	7.5.5	ePRR-r:M	ePRR-r:M	
edCSR-s	<s> Complete source routing	7.5.4	eCSR:M	eCSR:M	
edPSR-s	<s> Partial source routing	7.5.4	ePSR:M	X	
edQOSM-s	<s> QoS maintenance	7.5.6	c1:M	M	
edQOSM-r	<r> QoS maintenance	7.5.6	c2:M	M	
edPri-s	<s> Priority	7.5.7	ePri-s:M	ePri-s:M	
edPri-r	<r> Priority	7.5.7	ePri-r:M	ePri-r:M	
edData-s	<s> Data	7.6	M	M	
edData-r	<r> Data	7.6	M	M	
edUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded and where appropriate an Error Report PDU generated?	6.21	M	M	
edUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries: c1: eQOSM-s OR eCong-s
c2: eQOSM-r OR eCong-r

C.4.3.4.2 ER Parameters

Item	Parameter	Reference	Status	Profile	Support
eeFxPt-s	<s> Fixed part	7.2	M	M	
eeFxPt-r	<r> Fixed part	7.2	M	M	
eeAddr-s	<s> Addresses	7.3	M	M	
eeAddr-r	<r> Addresses	7.3	M	M	
eePadd-s	<s> Padding	7.5.2	ePadd-s:M	ePadd-s:M	
eePadd-r	<r> Padding	7.5.2	M	M	
eeSecu-s	<s> Security	7.5.3	eSecu-s:M	eSecu-s:M	
eeSecu-r	<r> Security	7.5.3	eSecu-r:M	eSecu-r:M	
eeCRR-s	<s> Complete route recording	7.5.5	eCRR-s:M	eCRR-s:M	
eeCRR-r	<r> Complete route recording	7.5.5	eCRR-r:M	eCRR-r:M	
eePRR-s	<s> Partial route recording	7.5.5	ePRR-s:M	ePRR-s:M	
eePRR-r	<r> Partial route recording	7.5.5	ePRR-r:M	ePRR-r:M	
eeCSR-s	<s> Complete source routing	7.5.4	eCSR:M	eCSR:M	
eePSR-s	<s> Partial source routing	7.5.4	ePSR:M	X	
eeQOSM-s	<s> QoS maintenance	7.5.6	c1:M	M	
eeQOSM-r	<r> QoS maintenance	7.5.6	c2:M	M	
eePri-s	<s> Priority	7.5.7	ePri-s:M	ePri-s:M	
eePri-r	<r> Priority	7.5.7	ePri-r:M	ePri-r:M	
eeData-s	<s> Data	7.6	M	M	
eeData-r	<r> Data	7.6	M	M	
eeUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded?	6.21	M	M	
eeUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries: c1: eQOSM-s OR eCong-s
c2: eQOSM-r OR eCong-r

C.4.3.4.3 Inactive Network Layer Protocol PDU Parameters

Item	Parameter	Reference	Status	Profile	Support
eiNLPI-s	<s> Inactive network layer protocol identifier	7.8.2	IAS-s:M	X	
eiNLPI-r	<r> Inactive network layer protocol identifier	7.8.2	IAS-r:M	M(10)	
eiData-s	<s> Data	7.8.3	IAS-s:M	X	
eiData-r	<r> Data	7.8.3	IAS-r:M	M(10)	

Notes: 10 See note 3

C.4.3.5 Timers

Item	Timer	Reference	Status	Values	Profile/ Supported Range & Default	Support/ Values supported
eLifReas	Is reassembly timer <= received derived PDU lifetime?	6.8	M		M(11)	
eReasLim	What values of the reassembly timer are supported?	6.8		500 ms to 127.5 s	500 ms to 127.5 sec/ Default of 12 sec	

Notes: 11 This corrects a defect in Table 9 of T1.204-1993.

C.4.4 Intermediate Systems

C.4.4.1 Applicability

The profile items in Section C.4.4 are applicable only to intermediate system implementations; i.e., those in which item IS in Section C.4.2 is supported.

C.4.4.2 Supported Functions

Item	Function	Reference	Status	Profile	Support
iPDUC	PDU composition	6.1	M	M	
iPDUD	PDU decomposition	6.2	M	M	
iHFA	Header format analysis	6.3	M	M	
iPDUL	<s> PDU lifetime control	6.4	M	M	
iRout	Route PDU	6.5	M	M	
iForw	Forward PDU	6.6	M	M	
iSegm	Segment PDU	6.7	iDSNS:M	M	
iReas	Reassemble PDU	6.8	O	i	
iDisc	Discard PDU	6.9	M	M	
iErep	Error reporting	6.10	M	M	
iEdec	<s> Header error detection	6.11	M	M	
*iSecu	<s> Security	6.13	O	i	
*iCRR	<s> Complete route recording	6.15	O	i	
*iPRR	<s> Partial route recording	6.15	O	i	
*iCSR	Complete source routing	6.14	O	i	
*iPSR	Partial source routing	6.14	O	X	
*iPri	<s> Priority	6.17	O	i	
*iQOSM	<s> QoS maintenance	6.16	O	M	
*iCong	<s> Congestion notification	6.18	O	see note 12	
iPadd	<s> Padding	6.12	M	M	
iEreq	Echo request	6.19	O	i(13)	
iErsp	Echo response	6.20	O	i(13)	
iSegS	Create segments smaller than necessary	6.8	O	i	
iDSNS	Simultaneous support of subnetworks with different SN-User data sizes	Table 9 note 3	O	M	

Notes:

12 See note 6.

13 See note 7.

C.4.4.3 Supported PDUs

Item	NPDU	Reference	Status	Profile	Support
iDT-t	DT (full protocol) transmit	7.7	M	M	
iDT-r	DT (full protocol) receive	7.7	M	M	
iDTNS-t	DT (non-segmenting) transmit	7.7	M	X(14)	
iDTNS-r	DT (non-segmenting) receive	7.7	M	M	
iER-t	ER transmit	7.9	M	M	
iER-r	ER receive	7.9	M	M	
iERQ-t	ERQ transmit	7.10	iEreq:M	i(15)	
iERQ-r	ERQ receive	7.10	M	i(15)	
iERP-t	ERP transmit	7.11	iErs:M	i(15)	
iERP-r	ERP receive	7.11	M	i(15)	

- Notes: 14 Implementation of this function is mandatory in the base PICS, but its use is **prohibited** in this profile.
15 See note 7. Note that implementation of this function may cause a backwards incompatibility with 1988 version of ISO/IEC 8473.

C.4.4.4 Supported Parameters

C.4.4.4.1 DT Parameters

Item	Parameter	Reference	Status	Profile	Support
idFxPt-s	<s> Fixed part	7.2	M	M	
idFxPt-r	<r> Fixed part	7.2	M	M	
idAddr-s	<s> Addresses	7.3	M	M	
idAddr-r	<r> Addresses	7.3	M	M	
idSeg-s	<s> Segmentation part	7.4	M	M	
idSeg-r	<r> Segmentation part	7.4	M	M	
idPadd-s	<s> Padding	7.5.2	M	M	
idPadd-r	<r> Padding	7.5.2	M	M	
idSecu-s	<s> Security	7.5.3	iSecu:M	iSecu:M	
idSecu-r	<r> Security	7.5.3	iSecu:M	iSecu:M	
idCRR-s	<s> Complete route recording	7.5.5	iCRR:M	iCRR:M	
idCRR-r	<r> Complete route recording	7.5.5	iCRR:M	iCRR:M	
idPRR-s	<s> Partial route recording	7.5.5	M	M	
idPRR-r	<r> Partial route recording	7.5.5	iPRR:M	iPRR:M	
idCSR-s	<s> Complete source routing	7.5.4	iCSR:M	iCSR:M	
idCSR-r	<r> Complete source routing	7.5.4	iCSR:M	iCSR:M	
idPSR-s	<s> Partial source routing	7.5.4	M	X	
idPSR-r	<r> Partial source routing	7.5.4	iPSR:M	X	
idQOSM-s	<s> QoS maintenance	7.5.6	M	M	
idQOSM-r	<r> QoS maintenance	7.5.6	cl:M	M	
idPri-s	<s> Priority	7.5.7	M	M	
idPri-r	<r> Priority	7.5.7	iPri:M	iPri:M	
idData-s	<s> Data	7.6	M	M	
idData-r	<r> Data	7.6	M	M	
idUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded and where appropriate an Error Report PDU generated?	6.21	M	M	
idUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries: cl: iQOSM OR iCong

C.4.4.4.2 ER Parameters

Item	Parameter	Reference	Status	Profile	Support
ieFxFt-s	<s> Fixed part	7.2	M	M	
ieFxFt-r	<r> Fixed part	7.2	M	M	
ieAddr-s	<s> Addresses	7.3	M	M	
ieAddr-r	<r> Addresses	7.3	M	M	
iePadd-s	<s> Padding	7.5.2	M	M	
iePadd-r	<r> Padding	7.5.2	M	M	
ieSecu-s	<s> Security	7.5.3	iSecu:M	iSecu:M	
ieSecu-r	<r> Security	7.5.3	iSecu:M	iSecu:M	
ieCRR-s	<s> Complete route recording	7.5.5	iCRR:M	iCRR:M	
ieCRR-r	<r> Complete route recording	7.5.5	iCRR:M	iCRR:M	
iePRR-s	<s> Partial route recording	7.5.5	M	M	
iePRR-r	<r> Partial route recording	7.5.5	iPRR:M	iPRR:M	
ieCSR-s	<s> Complete source routing	7.5.4	iCSR:M	iCSR:M	
ieCSR-r	<r> Complete source routing	7.5.4	iCSR:M	iCSR:M	
iePSR-s	<s> Partial source routing	7.5.4	M	X	
iePSR-r	<r> Partial source routing	7.5.4	iPSR:M	X	
ieQOSM-s	<s> QoS maintenance	7.5.6	M	M	
ieQOSM-r	<r> QoS maintenance	7.5.6	cl:M	M	
iePri-s	<s> Priority	7.5.7	M	M	
iePri-r	<r> Priority	7.5.7	iPri:M	iPri:M	
ieData-s	<s> Data	7.6	M	M	
ieData-r	<r> Data	7.6	M	M	
ieUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded?	6.21	M	M	
ieUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries: cl: iQOSM OR iCong

C.4.4.5 Timer and Parameter Values

Item	Timer	Reference	Status	Values	Profile/ Supported Range & Default	Support/ Values supported
iLifReas	Is reassembly timer <= received derived PDU lifetime?	6.8	iReas: M		iReas:M	
iReasLim	What values of the reassembly timer are supported?	6.8		500 ms to 127.5 s	500 ms to 127.5 seconds/ Default of 12 seconds	

C.5 ISO 9542 SONET Protocol Profile

The following SONET profile specifies the implementation requirements for the End system to Intermediate system routing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473). The protocol profile is based on the PICS proforma contained in Annex A of ISO 9542: 1988, with an additional Bellcore profile column.

C.5.1 Notations

C.5.1.1 Status Symbols

M mandatory

O optional

X prohibited

CI: the status following this symbol applies only when the profile states that configuration information is supported.

RI: the status following this symbol applies only when the profile states that redirection information supported.

(CI∨RI): the status following this symbol applies only when the profile states that either configuration information or redirection information (or both) is supported.

C.5.1.2 Other Symbols

<r> receive aspects of an item

<s> send aspects of an item

C.5.2 PICS Proforma: ISO 9542(1988) – End System

Item	Protocol Function	References	Status	Profile	Support
CI	Is configuration information supported?		O	ENE:M:1	
RI	Is redirection information supported?		O	ENE:M:2	

Notes:

- 1 Support is mandatory for both LAN and DCC operations communications.
- 2 Support is mandatory for LAN operations communications only.

C.5.2.1 Supported Functions

Item	Function	References	Status	Profile	Support
CfRs	Configuration Response	6.6	M	M	
ErrP	Protocol Error Processing	6.13	(CI∨RI): M	M	
HCsV	PDU Header Checksum Validation	6.12	(CI∨RI): M	M	
HCsG	PDU Header Checksum Generation	6.12	O	i	
RpCf	Report Configuration	6.2, 6.2.1	CI:M	M	
RcCf	Record Configuration	6.3, 6.3.2	CI:M	M	
FlCf	Flush Old Configuration	6.4	CI:M	M	
QyCf	Query Configuration	6.5	CI:M	M	
RcRd	Record Redirect	6.9	RI:M	RI:M	
FlRd	Flush Old Redirect	6.11	RI:M	RI:M	
RfRd	Refresh Redirect	6.10	RI:O	RI:M	
CfNt	Configuration Notification	6.7	CI:O	i	
CTPr	ESCT Processing	6.3.2	CI:O	M	
AMPr	Address Mask (only) Processing	7.4.5	RI:O	see note 3	
SMPr	Address Mask and SNPA Mask Processing	7.4.5, 7.4.6	RI:O	see note 3	

Notes:

- 3 This is a **Bellcore objective**. It is defined to be **out of scope** for this profile. This is different than T1.204-1993.

C.5.2.2 Supported PDUs

Item	NPDU	References	Status	Profile	Support
ESH-s	<s> End System Hello	7.1, 7.5	M	M	
ESH-r	<r> End System Hello	7.1, 7.5	CI:M	M	
ISH-r	<r> Intermediate System Hello	7.1, 7.5	CI:M	M	
RD-r	<r> Redirect	7.1, 7.5	RI:M	RI:M	

C.5.2.3 Supported Parameters

Item	Parameter	References	Status	Profile	Support
FxPt-s	<s> Fixed Part	7.2.1-7.2.7	M	M	
FxPt-r	<r> Fixed Part	7.2.1-7.2.7	(CIvRI): M	M	
SA-sl	<s> Source Address, one NSAP only	7.3.1, 7.3.2	O:I	M	
SA-rl	<r> Source Address, one NSAP only	7.3.1, 7.3.2	CI:M	M	
SA-sm	<s> Source Address, two or more NSAPs	7.3.3	O:I	M	
SA-rm	<r> Source Address, two or more NSAPs	7.3.3	CI:M	M	
NET-r	<r> Network Entity Title	7.3.1, 7.3.2, 7.3.4	(CIvRI): M	M	
DA-r	<r> Destination Address	7.3.1, 7.3.2, 7.3.5	RI:M	RI:M	
BSNPA-r	<r> Subnetwork Address	7.3.1, 7.3.2, 7.3.6	RI:M	RI:M	
Scty-s	<s> Security	7.4.2	O	see note 4	
Scty-r	<r> Security	7.4.2	O	see note 4	
Pty-s	<s> Priority	7.4.4	O	i	
Pty-r	<r> Priority	7.4.4	O	i	
QoS-r	<r> QoS Maintenance	7.4.3	RI:O	i	
AdMk-r	<r> Address Mask	7.4.5	RI:O	see note 5	
SNMK-r	<r> SNPA Mask	7.4.6	RI:O	see note 5	
ESCT-r	<r> Suggested ES Configuration Timer	7.4.7	CI:O	M	
OOpt-r	<r> (ignore) unsupported or unknown options		M	M	
OOpt-s	<s> Other options	7.4.1	X	X	

Notes:

- 4 This is a **Belcore objective**. It is defined to be **out of scope** for this profile.
5 See note 3.

C.5.2.4 Supported Parameter Ranges

Item	Ranges	References	Status	Profile/ Supported Range & Default	Support
HTv	What range of values can be set for the Holding Time field in transmitted PDUs?	6.1, 6.1.2	M	M/ 1 sec to 500 sec & default of 105 seconds	
CTv	If configuration information is supported, what range of values can be set for the Configuration Timer?	6.1, 6.1.1	CI:M	M/ 1 sec to 200 sec & default of 50 seconds	

C.5.3 PICS Proforma: ISO 9542(1988) – Intermediate System

Item	Protocol Function	References	Status	Profile	Support
CI	Is configuration information supported?		O	INE:M:6 GNE:M:6	
RI	Is redirection information supported?		O	INE:M:7 GNE:M:7	

Notes:

- 6 See note 1.
- 7 See note 2.

C.5.3.1 Supported Functions

Item	Protocol Function	References	Status	Profile	Support
ErrP	Protocol Error Processing	6.13	M	M	
HCsV	PDU Header Checksum Validation	6.12	M	M	
HCsG	PDU Header Checksum Generation	6.12	O	i	
RpCf	Report Configuration	6.2, 6.2.2	CI:M	M	
RpCf	Record Configuration	6.3, 6.3.1	CI:M	M	
FICf	Flush Old Configuration	6.4	CI:M	M	
RqRd	Request Redirect	6.8	RI:M	RI:M	
CfNt	Configuration Notification	6.7	CI:O	i	
CTGn	ESCT Generation	6.3.2	CI:O	M	
AMGn	Address Mask (only) Generation	6.8	RI:O	see note 8	
SMGn	Address Mask and SNPA Mask Generation	6.8	RI:O	see note 8	

Notes:

- 8 See note 3.

C.5.3.2 Supported PDUs

Item	NPDU	References	Status	Profile	Support
ESH-r	<s> End System Hello	7.1, 7.5	CI:M	M	
ISH-r	<r> Intermediate System Hello	7.1, 7.6	CI:O	M	
ISH-s	<s> Intermediate System Hello	7.1, 7.6	CI:M	M	
RD-s	<s> Redirect	7.1, 7.7	RI:M	RI:M	
RD-r	<r> Redirect	6.9, 7.1, 7.7	M	M	

C.5.3.3 Supported Parameters

Item	Parameter	References	Status	Profile	Support
FxPt-s	<s> Fixed Part	7.2.1-7.2.7	M	M	
FxPt-r	<r> Fixed Part	7.2.1-7.2.7	M	M	
SA-r	<s> Source Address, two or more NSAPs	7.3.1,7.3.2, 7.3.3	CI:M	M	
NET-s	<s> Network Entity Title	7.3.1,7.3.2, 7.3.4	M	M	
NET-r	<r> Network Entity Title	7.3.1,7.3.2, 7.3.4	ISH-r:M	M	
DA-s	<s> Destination Address	7.3.1,7.3.2, 7.3.5	RI:M	RI:M	
BSNPA-s	<s> Subnetwork Address	7.3.1,7.3.2, 7.3.6	RI:M	RI:M	
Scty-s	<s> Security	7.4.2	O	see note 9	
Scty-r	<r> Security	7.4.2	O	see note 9	
Pty-s	<s> Priority	7.4.4	O	i	
Pty-r	<r> Priority	7.4.4	O	i	
QoS M-s	<s> QoS Maintenance	7.4.3	RI:O	i	
AdMk-s	<s> Address Mask	7.4.5	RI:O	see note 10	
SNMK-s	<s> SNPA Mask	7.4.6	RI:O	see note 10	
ESCT-s	<s> Suggested ES Configuration Timer	7.4.7	CI:O	M	
ESCT-r	<r>(ignore)Suggested ES Configuration Timer	7.4.7	ISH-r:M	M	
OOpt-r	<r> (ignore) unsupported or unknown options		M	M	
OOpt-s	<s> Other options	7.4.1	X	X	

Notes:

- 9 See note 4.
10 See note 3.

C.5.4 Supported Parameter Ranges

Item	Ranges	References	Status	Profile/ Supported Range & Default	Support
HTv	What range of values can be set for the Holding Time field in transmitted PDUs?	6.1, 6.1.2	M	M/ 1 sec to 500 sec & default of 25 seconds	
CTv	If configuration information is supported, what range of values can be set for the Configuration Timer?	6.1, 6.1.1	CI:M	M/ 1 sec to 200 sec & default of 10 seconds	

C.6 ISO 10589 Protocol Profile

The following profile specifies the implementation requirements for the Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol providing the connectionless-mode network service (ISO 8473). The protocol profile is based on the PICS proforma contained in Annex A of ISO 10589: 1992, with an additional Bellcore profile column.

C.6.1 Notations (Status Symbols)

- M mandatory
- O optional
- O.<n> optional, but support of at least one of the group of options labeled by the same numeral <n> is required
- X prohibited
- c.<n> conditional requirement, according to condition <p>
- not applicable
- * Items whose references are used in predicates are indicated by an asterisk in the Item column.

C.6.2 Protocol Summary: ISO 10589 General

Item	Functionality/Description	References	Status	Profile	Support
AllIS	Are all basic IS-IS routing functions implemented?	12.1.2	M	M	
System Management	Is the system capable of being managed by the specified management information?	11	M	M(3)	
Authentication	Is PDU authentication based on passwords implemented?	7.3.7-7.3.10, 7.3.15.1-7.3.15.4, 8.2.3-8.2.4, 8.4.1.1	O	see note 1	
Default Metric	Is the default metric supported?	7.2.2, 7.2.6	M	M	
Delay Metric	Is the delay metric supported?	7.2.2, 7.2.6	O	i	
Expense Metric	Is the expense metric supported?	7.2.2, 7.2.6	O	i	
Error Metric	Is the error metric supported?	7.2.2, 7.2.6	O	i	
ID Field Length	What values of RouteingDomainIDLength (from the set 1-8) are supported by this implementation. Is the value configurable by system management?	7.1.3	M	M(2) NO	
Forwarding Rate	How many ISO 8473 PDUs can the implementation forward per second?	12.2.5.1.b	M	M	
Performance	Are the implementation performance criteria met?	12.2.5	M	M	

- Notes:
- 1 This is a **Bellcore objective**. It is defined to be **out of scope** for this profile.
 - 2 The value supported is 6 octets.
 - 3 If TL1 is used as the application layer protocol for management, then the NE shall provide management functionality equivalent to that which CMISE provides.

C.6.2.1 System Environment: General

Item	Functionality/Description	References	Status	Profile	Support
ISO9542	Are the appropriate ISO 9542 operations implemented?	10.3, 8.2.1-8.2.2, 8.3.4, 8.4.5, 8.4.6	M	M	
Timer Jitter	Is jitter introduced in all periodic timers whose expiration causes transmission of a PDU?	10.1	M	M	

C.6.2.2 Subnetwork Dependent Functions: General

Item	Functionality/Description	References	Status	Profile	Support
*LAN	Are the subnetwork dependent functions for broadcast subnetworks implemented?	8.4	O.1	M:4	
LAN IS Adjacencies	Are the LAN IS adjacency establishment operations implemented?	8.4.1-8.4.3	LAN: M	LAN: M	
LAN ES Adjacencies	Are the LAN ES adjacency establishment operations implemented?	8.4.6	LAN: M	LAN: M	
LAN DIS	Are the LAN designated IS operations implemented?	8.4.4, 8.4.5	LAN: M	LAN: M	
*8208 Static	Are the subnetwork dependent functions for ISO 8208 subnetworks implemented?	8.3	O.1	i	
8208 SNDCF	Are the ISO 8208 Subnetwork Dependent Convergence Functions implemented?	8.3.1, 8.3.2.1	C.1:M	i	
*PtPt	Are the subnetwork dependent functions for point-to-point subnetworks implemented?	8.2	O.1	M:5	
PtPt IS Adjacencies	Are the point-to-point IS adjacency establishment operations implemented?	8.2.2-8.2.5	C.2:M	PtPt:M	
PtPt ES Adjacencies	Are the point-to-point ES adjacency establishment operations implemented?	8.2.1	C.2:M	PtPt:M	
PtPt IIH PDU	Are point-to-point IIH PDUs correctly constructed and parsed?	9.7	C.2:M	PtPt:M	

C.1 if 8208 Static or 8208 DA then M else -

C.2 if PtPt or 8208 Static then M else -

Notes:

- 4 Support is **mandatory** for LAN.
- 5 Support is **mandatory** for DCC.

C.6.2.3 Update Process: General

Item	Functionality/Description	References	Status	Profile	Support
LSP Periodic Generation	Is periodic generation of new local LSPs implemented?	7.3.2, 7.3.5, 7.3.13	M	M	
LSP Event Driven Generation	Is event driven generation of new local LSPs implemented?	7.3.6	M	M	
Pseudonode LSP Generation	Is generation of pseudonode LSPs implemented?	7.3.8, 7.3.10	LAN: M	LAN: M	
Multiple LSP Generation	Is multiple LSP generation implemented?	7.3.4	M	M	
LSP Propagation	Is propagation of LSPs implemented?	7.3.12, 7.3.14, 7.3.15.1, 7.3.15.5	M	M	
LSP Lifetime Control	Are the LSP lifetime control operations implemented?	7.3.16.4, 7.3.16.3	M	M	
CSNP Generation	Is the generation of CSNPs implemented?	7.3.15.3, 7.3.17	M	M	
PSNP Generation	Is the generation of PSNPs implemented?	7.3.15.4, 7.3.17	M	M	
SNP Processing	Are the sequence number PDU processing procedures implemented?	7.3.15.2, 7.3.17	M	M	
LSDB Overload	Are the LSP database overload operations implemented?	7.3.19	M	M	

C.6.2.4 Decision Process: General

Item	Functionality/Description	References	Status	Profile	Support
Minimum Cost Path	Is computation of a single minimum cost path based upon each supported metric implemented?	7.2.6	M	M	
Equal Cost Paths	Is computation of equal minimum cost paths based upon each supported metric implemented?	7.2.6	O	i	
Downstream Paths	Is computation of downstream routes based upon each supported metric implemented?	7.2.6	O	i	
Multiple LSPs Recognition	Are multiple LSPs used only when a LSP with LSP#0 and remaining lifetime greater than 0 is present?	7.2.5	M	M	
Overloaded IS Exclusion	Are links to ISs with overloaded LSDBs ignored?	7.2.8.1	M	M	
Two Way Connectivity	Are links not reported by both end ISs ignored?	7.2.8.2	M	M	
Path Preference	Is the order of preference for path selection implemented?	7.2.12	M	M	
Excess Path Removal	Is removal of excess paths implemented?	7.2.7	M	M	
FIB Construction	Is the construction of ISO8473 Forwarding Information Bases implemented?	7.2.9	M	M	

C.6.2.5 Forward/Receive Process: General

Item	Functionality/Description	References	Status	Profile	Support
FIB Selection	Is selection of appropriate Forwarding Information Base implemented?	7.4.2	M	M	
NPDU Forwarding	Is forwarding of ISO8473 PDUs implemented?	7.4.3.1, 7.4.3.3	M	M	
Receive Process	Are the basic receive process functions implemented?	7.4.4	M	M	

C.6.3 Protocol Summary: ISO 10589 Level 1 Specific Functions

Item	Functionality/ Description	References	Status	Profile/ Supported Range & Default	Support
*LIIS	Are Level 1 IS-IS routing functions implemented?	12.1.3	M	M	
Maximum Area Addresses	What values of maximumAreaAddresses are supported by this implementation?	7.1.5, 7.2.11	LIIS: M	M/ 0 to 12 & default of 3	
Area IS Count	How many ISs can this system support in a single area?	12.2.5	LIIS: M	M/ 1 to 512 & default of 512	
L1 Manual ES Adjacency	Are the manual ES adjacencies implemented?	7.3.31	LIIS: M	M	

C.6.3.1 Level 1 Subnetwork Dependent Functions

Item	Functionality/Description	References	Status	Profile	Support
L1 LAN IIH PDU	Are L1 LAN IIH PDUs correctly constructed and parsed?	9.5	C.3:M	C.3:M	

C.3 if LIIS and LAN then M else -

C.6.3.2 Level 1 Update Process

Item	Functionality/Description	References	Status	Profile	Support
L1 LS PDU	Are L1 LS PDUs correctly constructed and parsed?	9.8	LIIS: M	M	
L1 CSN PDU	Are L1 CSN PDUs correctly constructed and parsed?	9.10	LIIS: M	M	
L1 PSN PDU	Are L1 PSN PDUs correctly constructed and parsed?	9.12	LIIS: M	M	

C.6.3.3 Level 1 Decision Process

Item	Functionality/Description	References	Status	Profile	Support
L1 Nearest L2 IS Identification	Is the identification of the nearest L2 IS implemented?	7.2.9.1	LIIS: M	M	
L1 Area Addresses Computation	Is the computation of area addresses implemented?	7.2.11	LIIS: M	M	

C.6.4 Protocol Summary: ISO 10589 Level 2 Specific Functions

Item	Functionality/Description	References	Status	Profile/Supported Range & Default	Support
*L2IS	Are level 2 IS-IS routing functions implemented?	12.1.4	O	M:6	
IS Count	What is the total number of ISs that this L2 IS can support?	12.2.5	L2IS:M	L2IS:M/ 1 to 512 & default of 512 (see note 7)	
L2IS Count	How many level 2 ISs does this implementation support?	12.2.5.1	L2IS:M	L2IS:M/ 1 to 512 & default of 256 (see note 7)	
*RA Prefix	Are Reachable Address Prefixes supported on circuits?	8.1, 7.3.3.2	L2IS:O	L2IS:M	
External Metrics	Are external metrics supported?	7.2.2, 7.2.12, 7.3.3.2	RA Prefix:M	RA Prefix:M	
*Partition	Is level 1 partition repair implemented?	7.2.10	L2IS:O	see note 8	

Notes:

- 6 This function is **mandatory** when the Level 2 functions are supported.
- 7 Note that these numbers are preliminary and are subject to further study and possible change.
- 8 This is a **Bellcore objective**. It is defined to be **out of scope** for this profile.

C.6.4.1 Level 2 Subnetwork Dependent Functions

Item	Functionality/Description	References	Status	Profile	Support
L2 LAN IIH PDU	Are L2 LAN IIH PDUs correctly constructed and parsed?	9.6	C.4:M	C.4:M	
*8208 DA	Are ISO8208 Dynamic Assignment circuits implemented?	8.3	O.1	i	
RA Adjacency Management	Are the reachable address adjacency management operations implemented?	8.3.2.2-8.3.5.6	8208 DA:M	8208 DA:M	
Call Establishment Metric Increment	Are non-zero values of the callEstablishment-MetricIncrement supported?	8.3.5	8208 DA:O	i	
Reverse Path Cache	Is 8208 reverse path cache implemented?	8.3.3	8208 DA:O	i	

C.4 if L2IS and LAN then M else -

C.6.4.2 Level 2 Update Process

Item	Functionality/Description	References	Status	Profile	Support
L2 LS PDU	Are L2 LS PDUs correctly constructed and parsed?	9.9	L2IS: M	L2IS: M	
L2 CSN PDU	Are L2 CSN PDUs correctly constructed and parsed?	9.11	L2IS: M	L2IS: M	
L2 PSN PDU	Are L2 PSN PDUs correctly constructed and parsed?	9.13	L2IS: M	L2IS: M	

C.6.4.3 Level 2 Decision Process

Item	Functionality/Description	References	Status	Profile	Support
L2 Attached Flag	Is the setting of the attached flag implemented?	7.2.9.2	L2IS:M	L2IS:M	
L2 Partition DIS election	Is the election of partition L2 DIS implemented?	7.2.10.2	Partition: M	Partition: M	
L2 Partition Area Addresses Computation	Is the computation of L1 partition area addresses implemented?	7.2.10.3	Partition: M	Partition: M	
L2 DIS Partition Repair	Is partition detection and repair via virtual L1 links implemented?	7.2.10.1	Partition: M	Partition: M	

C.6.4.4 Level 2 Forward/Receive Process

Item	Functionality/Description	References	Status	Profile	Support
L2 NPDU Encapsulation	Is the encapsulation of NPDUs implemented?	7.2.10.4, 7.4.3.2	Partition: M	Partition: M	
L2 NPDU Decapsulation	Is the decapsulation of NPDUs implemented?	7.4.4	Partition: M	Partition: M	

C.7 ISO 8073 Protocol Profile

The following profile specifies the implementation requirements for the Connection-mode Transport Protocol. The protocol profile is based on the PICS proforma contained in Annex C of ISO 8073-1:1992, with an additional Bellcore profile column.

C.7.1 Notations

C.7.1.1 Status Symbols

M mandatory

O optional to implement. If implemented, the feature may or may not be used.

O.<n> optional, but support of at least one of the group of options labeled by the same numeral <n> is required.

<index>: this predicate symbol means that the status following it applies only when the profile states that the feature identified by the index is supported. In the simplest case, <index> is the identifying tag of a single profile item. <index> may also be a Boolean expression composed of several indices.

<index>:: when this group predicate is true, the associated clause should be completed.

C.7.2 Protocol Implementation for TP4/CLNS (C4L::)

C.7.2.1 Annex B – NCMS

Index	Class	References	Status	Profile	Support
A1	Network connection management procedures	Annex B	O	i	

C.7.2.2 Classes Implemented

Index	Class	References	Status	Profile	Support
C4L	Class 4 operation over CLNS	14	O	M	

C.7.3 Initiator/Responder Capability for Protocol Classes 0-4

Index	Item	References	Status	Profile	Support
IR1	Initiating CR TPDU	14.5 a)	O.2	M	
IR2	Responding to CR TPDU	14.5 a)	O.2	M	

C.7.4 Supported Functions

C.7.4.1 Supported Functions for Class 4 (C4L::)

The following functions are mandatory.

Index	Function	References	Status	Profile	Support
T4F1	TPDU transfer	6.2	M	M	
T4F2	Segmenting	6.3	M	M	
T4F3	Reassembling	6.3	M	M	
T4F4	Separation	6.4	M	M	
T4F5	Connection establishment	6.5	M	M	
T4F6	Connection refusal	6.6	M	M	
T4F7	Data TPDU numbering (normal)	6.10	M	M	
T4F8	Retention and acknowledgment of TPDUs Retention until acknowledgment of TPDUs (AK)	6.13.4.1	M	M	
T4F9	Explicit flow control	6.16	M	M	
T4F10	Checksum	6.17	M	M:1	
T4F11	Frozen references	6.18	M	M	
T4F12	Retransmission on time-out	6.19	M	M-	
T4F13	Resequencing	6.20	M	M	
T4F14	Inactivity control	6.21	M	M	

Notes:

- 1 Checksum is mandatory for CR TPDU only.

The following functions are mandatory if class 4 is operated over CLNS.

Index	Function	References	Status	Profile	Support
T4F23	Transmission over CLNS	6.1.2	M	M	
T4F24	Normal release when operating over CLNS (explicit)	6.7.2	M	M	
T4F25	Association of TPDUs with Transport connection when operating over CLNS	6.9.2	M	M	
T4F26	Expedited data transfer when operating over CLNS (Network normal)	6.11.2	M	M:2	
T4F27	Treatment of protocol errors when operating over CLNS	6.22.2	M	M	

Notes:

- 2 The support of this function is mandatory in ISO 8073. T1.204 calls for support of the expedited data transfer service with non-use being the default negotiation. Therefore, in generating CR and CC TPDUs, the transport layer entity shall set bit 1 of the Additional Option Selection optional parameter to 0 (zero) to negotiate non-use of the transport expedited data transfer service.

The following functions are optional.

Index	Function	References	Status	Profile	Support
T4F28	Data TPDU numbering (extended)	6.10	O	M:3	
T4F29	Non-use of checksum	6.17	O	M:4	
T4F30	Concatenation	6.4	O	i	
T4F31	Retention and acknowledgment of TPDUs - Use of selective acknowledgment	6.13.4.4	O	i	
T4F32	Retention and acknowledgment of TPDUs - Use of request acknowledgment	6.13.4.3	O	i	

Notes:

- 3 EXTENDED must be supported, but NORMAL is the default setting.
- 4 All TPDUs, except CR TPDU, shall negotiate "non-use" of the checksum. Initiators shall request and responders shall agree to "non-use" of the checksum.

C.7.5 Supported TPDUs

The following TPDUs and the parameters that constitute their fixed parts are mandatory if a corresponding predicate in the status column is true.

Index	TPDUs		References	Status	Profile	Support
ST1	CR	supported on transmission	13.1	IR1:M	M	
ST2	CR	supported on receipt	13.1	IR2:M	M	
ST3	CC	supported on transmission	13.1	IR2:M	M	
ST4	CC	supported on receipt	13.1	IR1:M	M	
ST5	DR	supported on transmission	13.1	IR2:M	M	
ST6	DR	supported on receipt	13.1	IR1:M	M	
ST7	DC	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST8	DC	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST9	DT	supported on transmission	13.1	M	M	
ST10	DT	supported on receipt	13.1	M	M	
ST11	ED	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST12	ED	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST13	AK	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST14	AK	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST15	EA	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST16	EA	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST19	ER	supported on receipt	13.1	M	M	

Notes:

C1	Class 1
C2	Class 2
C3	Class 3
C4	Class 4 over CONS

State for which classes, if any, ER is supported on transmission

Index	Class	References	Status	Profile	Support
SER4L	Class 4 over CLNS	6.22.2	O	M(5)	

Notes: 5 See special cases listed in the GR for handling certain errors in CR and CC TPDUs.

C.7.6 Supported Parameters of Issued TPDU

C.7.6.1 Parameter Values for CR TPDU (C4L::)

If the additional options selection parameter is issued in a CR TPDU it is mandatory that

Index		References	Profile	Support
ICR1	Bits 8 and 7 shall be set to zero	13.3.4 g)	M	

C.7.6.2 Supported Parameters for Class 4 TPDU (C4L::)

The following parameters are optional if a CR TPDU is issued with preferred class 4

Index	Supported parameter	References	Status	Profile	Support
I4CR7	Called TSAP-ID	13.3.4 a)	O	M(6)	
I4CR8	Calling TSAP-ID	13.3.4 a)	O	M(6)	
I4CR9	TPDU size	13.3.4 b)	O	M	
I4CR10	Version number	13.3.4 d)	O	i	
I4CR11	Protection parameters	13.3.4 e)	O	i	
I4CR12	Additional option selection	13.3.4 g)	O	M(7)	
I4CR13	Throughput	13.3.4 k)	O	i	
I4CR14	Residual error rate	13.3.4 m)	O	i	
I4CR15	Priority	13.3.4 n)	O	i	
I4CR16	Transit delay	13.3.4 p)	O	i	
I4CR17	Acknowledge time	13.3.4 j)	O	i	
I4CR18	Preferred maximum TPDU size	13.3.4 c)	O	i(8)	
I4CR19	Inactivity time	13.3.4 r)	O	i(8)	

Notes:

- 6 The value of the Called TSAP-ID and Calling TSAP-ID parameters shall be ASCII "TT" (i.e., "5454" hex) to indicate that the ISO Session Layer is being run over TP4.
- 7 All TPDU, except CR TPDU, shall negotiate "non-use" of the checksum. Initiators shall request "non-use" of the checksum. See also Note 2.
- 8 This is a new optional function introduced in the 1992 version of ISO/IEC 8073. To be compatible with ISO/IEC ISP 10608-1 (which is based on ISO/IEC 8073:1988/Amd 3 1992), this new functionality is defined to be **out of scope** for this profile.

The following parameters are optional if a CC TPDU is issued in class 4.

Index	Supported parameters	References	Status	Profile	Support
I4CC6	Called TSAP-ID	13.4.4	O	M(9)	
I4CC7	Calling TSAP-ID	13.4.4	O	M(9)	
I4CC8	TPDU size	13.4.4	O	M	
I4CC9	Protection parameters	13.4.4	O	i	
I4CC10	Additional option selection	13.4.4	O	M(10)	
I4CC11	Acknowledge time	13.4.4	O	i	
I4CC12	Throughput	13.4.4	O	i	
I4CC13	Residual error rate	13.4.4	O	i	
I4CC14	Priority	13.4.4	O	i	
I4CC15	Transit delay	13.4.4	O	i	
I4CC16	Preferred maximum TPDU size	13.4.4	O	i(11)	
I4CC17	Inactivity time	13.4.4	O	i(11)	

- Notes: 9 See Note 6.
10 All TPDU's, except CR TPDU, shall negotiate "non-use" of the checksum. Responders shall agree to "non-use" of the checksum. See also Note 2.
11 See note 8.

The following parameter is optional if a DR TPDU is issued in class 4.

Index	Supported parameter	References	Status	Profile	Support
I4DR4	Additional information	13.5.4 a)	O	i	

The following parameter is mandatory in a DT TPDU if request of acknowledgment has been selected.

Index	Supported parameter	References	Status	Profile	Support
I4DT4	ROA	13.7.3 a)	O	i(12)	

- Notes: 12 See note 7.

The following parameter is mandatory in an AK TPDU if issued in class 4.

Index	Supported parameter	References	Status	Profile	Support
I4AK4	Flow control confirmation	13.9.4 c)	O	M	

If the implementation can reduce credit and does so in the manner outlined in 12.2.3.8.2 then subsequence number in AK TPDU's is mandatory. Otherwise complete item I4AK5.

Index	Supported parameter	References	Status	Profile	Support
I4AK5	Subsequence number	13.9.4 b)	O	M	

The following parameter is optional in an AK TPDU if selective acknowledgment has been negotiated.

Index	Supported parameter	References	Status	Profile	Support
I4AK6	Selective acknowledgment parameters	13.9.4 d)	O	i(13)	

Notes: 13 See note 8.

The following parameter is optional if a ER TPDU is issued in class 4.

Index	Supported parameter	References	Status	Profile	Support
I4ER3	Invalid TPDU	13.12.4 a)	O	i	

C.7.7 Supported Parameters for Received TPDUs

Implementors should be aware that implementations shall be capable of receiving and processing all possible parameters for all possible TPDUs, dependent upon the class and optional functions implemented.

C.7.8 User Data in Issued TPDUs

A TS-user may issue data with a T-CONNECT request, T-CONNECT response or T-DISCONNECT request. Then it shall be possible to send user data as follows:

C.7.8.1 Class 4 (C4L::)

Index	User data	References	Status	Profile	Support
D4ICR	User data of up to 32 octets in a CR with preferred class 4	13.3.5	O	X(14)	
D4ICC	User data of up to 32 octets in a CC	13.4.5	O	X(14)	
D4IDR	User data of up to 64 octets in a DR	13.5.5	O	i	

Notes: 14 Implementation of this function is optional in the base PICS, but its use is **prohibited** in the profile. No protocol implementations shall send user data in CR and CC TPDUs. See discussion in T1.204-1993.

C.7.9 User Data in Received TPDUs

For classes 1 to 4, if it is possible to initiate a CR TPDU then it shall be possible to receive the following.

Index	User data	References	Profile	Support
DRCC	32 octets of user data in a CC TPDU	13.4.5	M(15)	
DRDR	64 octets of user data in a DR TPDU	13.5.5	M	

Notes:

- 15 All protocol implementations shall be prepared to receive user data in CC TPDUs, and all implementations may ignore user data, i.e., user data shall not cause a disconnect.

For classes 1 to 4, if it is possible to respond to a CR TPDU then it shall be possible to receive the following.

Index	User data	References	Profile	Support
DRCR	32 octets of user data in a CR TPDU	13.3.5	M(16)	

Notes:

- 16 All protocol implementations shall be prepared to receive user data in CR TPDUs, and all implementations may ignore user data, i.e., user data shall not cause a disconnect.

C.7.10 Negotiation

C.7.10.1 Class Negotiation - Initiator

What class(es) is (are) contained in the alternative class parameter if the preferred class is:

Index	Preferred class	References	Allowed values	Profile values	Support
NAC5	Class 4 over CLNS	6.5.5 j)	None	None	

C.7.10.2 TPDU Size Negotiation

Index		References	Status	Profile	Support
TS1	If maximum TPDU size is proposed in a CR TPDU then the initiator shall support all TPDU sizes from 128 octets to the maximum proposed.	14.6	M	M	
TS2	If the preferred maximum TPDU size parameter is used in a CR TPDU then the initiator shall support all TPDU sizes, except 0, that are multiples of 128 octets up to the preferred maximum proposed.	14.6 e)	I4CR18: M	i(17)	

Notes: 17 Implementation of this function is conditionally mandatory in the base PICS. To be compatible with ISO/IEC ISP 10608-1 (which is based on ISO/IEC 8073:1988/Amd 3 1992), use of this functionality is defined to be **out of scope** for this profile.

Index	TPDU size	References	Allowed values	Profile	Support
T4S1	What is the largest value of the maximum TPDU size parameter in a CR TPDU with preferred class 4?	14.6 e)	NOT I4CR18: One of 128, 256, 512, 1024, 2048, 4096, 8192 I4CR18: One of n128 with n = 1, 2, 3,...	1024, see note 18	
T4S2	What is the largest value of the maximum TPDU size parameter which may be sent in a CC TPDU when class 4 is selected?	14.6 e)	NOT I4CC16: One of 128, 256, 512, 1024, 2048, 4096, 8192 I4CC16: One of n128 with n = 1, 2, 3, ...	1024, see note 18	

Notes: 18 Note larger TPDU sizes (2048, 4096, and 8192) are optional. See T1.204-1993.

C.7.10.3 Use of Extended Format

Index	Extended format	References	Allowed values	Profile values	Support
NEF3	What formats can you propose in the CR TPDU in class 4?	6.5.5 n)	normal, extended	see note 19	
NEF6	What formats can you select in CC when extended has been proposed in CR class 4?	6.5.5 n)	normal, extended	see note 19	

Notes:

19 Extended format options shall be implemented. Non-use of extended format shall be negotiable. The responder shall honor the initiator's request whenever possible. Negotiation to other than what has been requested shall occur only under abnormal conditions: for example, severe congestion, as determined by the implementor. Initiators shall be prepared to operate in the mode confirmed by the responder. Normal is the default format. See T1.204-1993.

C.7.10.4 Expedited Data Transport Service

Index		References	Status	Profile	Support
TED1	Expedited data indication in CR and CC TPDU	6.5.5 r)	M	M	

C.7.10.5 Non-use of Checksum (C4L AND T4F29::)

Index	Non-use checksum	References	Allowed values	Profile values	Support
NUC1	What proposals can you make in the CR?	6.5.5 p)	non-use, use	non-use	
NUC2	What proposals can you make in CC when non-use of checksum has been proposed in CR?	6.5.5 p)	non-use, use	non-use	

C.7.10.6 Use of Selective Acknowledgment (*See note 20*)

Index	Selective acknowledgment	References	Allowed values	Profile values	Support
USA1	Is use of selective acknowledgment proposed in CR TPDUs?	6.5.5 s)	Yes, No	i	
USA2	If use of selective acknowledgment selected in a CC when it has been proposed in a CR?	6.5.5 s)	Yes, No	i	

Notes:

20 This is a new function in the 1992 base PICS. To be consistent with ISO/IEC ISP 10608-1 (which is based on ISO/IEC 8073:1988/Amd 3 1992), this functionality is defined to be **out of scope** for this profile.

C.7.10.7 Use of Request of Acknowledgment (*See note 21*)

Index	Request of acknowledgment	References	Allowed values	Profile	Support
ROA1	Is use of request of acknowledgment proposed in CR TPDUs?	6.5.5 t)	Yes, No	i	
ROA2	Is use of request of acknowledgment selected in a CC when it has been proposed in a CR?	6.5.5 t)	Yes, No	i	

Notes: 21 See note 20.

C.7.11 Error Handling

C.7.11.1 Action on Receipt of a Protocol Error

Index	Item	References	Allowed values	Profile	Support
PE4L	Class 4 over CLNS	6.22.2.3	C4L: ER, DR, Discard	ER, DR, Discard, see note 22	

Notes: 22 See note 5.

C.7.11.2 Actions on Receipt of an Invalid or Undefined Parameter in a CR TPDU

Index	Event	References	Status	Profile	Support
RR1	A parameter not defined in ISO 8073 shall be ignored	13.2.3	M	M	
RR2	An invalid value in the alternative protocol class parameter shall be treated as a protocol error	13.2.3	M	M	
RR3	An invalid value in the class and option parameter shall be treated as a protocol error	13.2.3	M	M	
RR4	On receipt of the additional option* selection parameter bits 8 to 5, and bits 4 to 1 if not meaningful for the proposed class shall be ignored.	13.3.4	M	M	
RR5	If non-use of explicit flow control is proposed and bit 1 of the additional option selection parameter equals 1, it shall be treated as a protocol error.	13.2.3	M	M	
RR6	On receipt of the class and option parameter bits 4 to 1 if not meaningful for the proposed class shall be ignored	13.3.3	M	M(23)	

Notes:

23 This entry is not contained in ISO/IEC 8073:1988, but is contained in ISO/IEC 8073:1992 and it is included in the ISO/IEC 8073: 1992 base PICS for clarity. There are no incompatibilities between the PICS due to this entry.

What action is supported on receipt of the following?

Index	Event	References	Allowed actions	Profile	Support
RR7	A parameter defined in ISO/IEC 8073 (other than those covered above) and have an invalid value	13.2.3	ignore, protocol error	ignore	

C.7.11.3 Actions on Receipt of an Invalid or Undefined Parameter in a TPDU other than a CR TPDU

The following actions are mandatory.

Index	Event	References	Status	Profile	Support
UI1	A parameter not defined in ISO/IEC 8073 shall be treated as a protocol error	13.2.3	M	M	
UI2	A parameter which has an invalid value as defined in ISO/IEC 8073 shall be treated as a protocol error	13.2.3	M	M	
UI3 (class 4 only)	A TPDU received with a checksum which does not satisfy the defined formula shall be discarded.	6.17.3	M	M	

C.7.12 Timers and Protocol Parameters

The following are mandatory if class 4 is supported.

Index		References	Status	Profile/ Supported Range & Default	Support
TA1	<i>T_I</i>	12.2.1	M	M/ .25 seconds to 64 seconds & default of 8 seconds	
TA2	<i>N</i>	12.2.1	M	M/ 2 to 15 & default of 2	
TA3	<i>I_L</i>	12.2.1	M	M/ 2 second to 512 seconds & default of 64 seconds	
TA4	<i>W</i>	12.2.1	M	M/ 1 second to 256 seconds & default of 16 seconds	
TA5	<i>L</i>	12.2.1	M	M/ 1 second to 256 seconds & default of 32 seconds	

Index		References	Status	Profile	Support
OT9	Does IUT support optional timer <i>TS2</i> when operating in class 4?	6.22.2.3	O	i	

C.8 TARP Protocol Implementation Conformance Statement

The following tables serve two purposes. First, when the "Profile" column is removed, they serve as the PICS for the TARP protocol defined in Section 8.7 of this GR. Second, when the "Profile" column is kept, they serve as the profile for the implementation requirements for the TARP protocol. The symbols used in the status column for this PICS are the same as those defined in Section C.4.1 for the ISO CLNP PICS.

C.8.1 Major Function

Index	Functionality/Description	Reference	Status	Profile	Support
MF1	Does the NE support TARP on the NE-NE interface according to the requirements in GR-253-CORE?	8.7	O	M:1	
MF2	Does the NE support the finding of the NET that matches a given TID?	8.7.4.1	M	M	
MF3	Does the NE support the finding of the TID that matches a given NET?	8.7.4.2	M	M	
MF4	Support of notification to other NEs of TID or protocol address changes	8.7.4.3	M	M	
MF5	Does the NE transmit TARP PDUs within ISO 8473 (CLNP) Data (DT) PDUs?	8.7.1	M	M	

1. If SONET NE supports TL1/OSI on the NE-NE interface (DCC or LAN)

C.8.2 Supported PDUs

The NE shall provide the function of a TARP processor that is capable of supporting:

Index	Functionality/Description	Reference	Status	Profile	Support
SP1-s	TARP Type 1 PDUs on transmit	8.7.2.4, 8.7.5.1	M(2)	M(2)	
SP1-r	TARP Type 1 PDUs on receive	8.7.2.4, 8.7.5.1, 8.7.5.6	M(3)	M(3)	
SP2-s	TARP Type 2 PDUs on transmit	8.7.2.4, 8.7.5.1, 8.7.5.2	M(2)	M(2)	
SP2-r	TARP Type 2 PDUs on receive	8.7.2.4, 8.7.5.1, 8.7.5.6	M(3)	M(3)	
SP3-s	TARP Type 3 PDUs on transmit	8.7.2.4, 8.7.5.3	M	M	
SP3-r	TARP Type 3 PDUs on receive	8.7.2.4, 8.7.5.6	M	M	
SP4-s	TARP Type 4 PDUs on transmit	8.7.2.4, 8.7.5.4	M	M	
SP4-r	TARP Type 4 PDUs on receive	8.7.2.4, 8.7.5.6	M	M	
SP5-s	TARP Type 5 PDUs on transmit	8.7.2.4, 8.7.5.5	M	M	
SP5-r	TARP Type 5 PDUs on receive	8.7.2.4, 8.7.5.6	M	M	

- The tar-tor field is optional.
- The contents of the tar-tor field shall be ignored. The PDUs shall be processed correctly regardless of the presence or absence of the tar-tor field.

C.8.3 Protocol Specifications

C.8.3.1 TARP PDU CLNP Specifications

The following table defines the TARP PDU fields carried in the fixed part of the CLNP DT PDU.

Index	Functionality/Description	Reference	Status	Profile	Support
FxPt1	Lifetime of the CLNP DT PDU set to a value of 25000 milliseconds	8.7.1	M	M	
FxPt2	Segmentation Permitted Flag set to a value of one	8.7.1	M	M	
FxPt3	Error Report Flag set to a value of zero	8.7.1	M	M	

C.8.3.2 TARP PDU Specifications

The following table defines the TARP PDU fields carried in the data part of the CLNP DT PDU.

Index	Functionality/Description	Reference	Status	Profile	Support
Data1	TARP lifetime (tar-lif)	8.7.2.1	M	M	
Data2	TARP sequence number (tar-seq)	8.7.2.2	M	M	
Data3	Protocol address type (tar-pro)	8.7.2.3	M	M	
Data4	URC and TARP type code (tar-tcd)	8.7.2.4	M	M	
Data5	TID target length (tar-tln)	8.7.2.5	M	M	
Data6	TID originator length (tar-oln)	8.7.2.6	M	M	
Data7	Protocol address length (tar-pln)	8.7.2.7	M	M	
Data8	TID of target (tar-ttg)	8.7.2.8	M	M	
Data9	TID of originator (tar-tor)	8.7.2.9	M	M	
Data10	Protocol address of originator (tar-por)	8.7.2.10	M	M	
Data11	URC bit ignored upon receipt of TARP PDU	8.7.2.4	M	M	

C.8.3.3 Protocol Timer Specifications

Index	Timer	Reference	Status	Profile/ Supported Range & default	Support
PTS1	Timer T1	8.7.4.1	M	M/ 0 to 3600 seconds & default of 15 seconds	
PTS2	Timer T2	8.7.4.1	M	M/ 0 to 3600 seconds & default of 25 seconds	
PTS3	Timer T3	8.7.4.1	M	M/ 0 to 3600 seconds & default of 40 seconds	
PTS4	Timer T4	8.7.4.1	M	M/ 0 to 3600 seconds & default of 20 seconds	

C.8.4 Major Capabilities

Index	Functionality/Description	Reference	Status	Profile	Support
MC1	Does the NE have the capability of designating one NET as the NET to be used for the mapping purposes for all TARP-related TID/NET mappings?	8.7	O	M:4	
MC2	Does the NE return a TARP Type 3 PDU with the tar-oln field equal to zero on receipt of a TARP Type 5 PDU for an NET other than the "designated NET"?	8.7	MC1:M	M	
MC3	Support of a TARP Data Cache (TDC)	8.7.3	O	i	
MC4	Support the incrementing of the tar-seq field on the origination of a TARP PDU	8.7.5	M	M	
MC5	Generation of a TARP Type 4 PDU to ensure other NEs LDB's are reset to zero (i.e., broadcast Type 4 PDU with tar-seq=0 to all adjacencies)	8.7.5	O	i	
MC6	Support ES TARP PDU processing on receipt	8.7.5.6.1	O.1	M:5	
MC7	Support Level 1 IS TARP PDU processing on receipt	8.7.5.6.2, 8.7.5.8	O.1	M:6	
MC8	Support Level 2 IS TARP PDU processing on receipt	8.7.5.6.3, 8.7.5.8	O.1	M:7	
MC9	Support propagation procedures	8.7.5.8	M:8	M:8	
MC10	Support of a circular (first-in first-out)TARP Loop Detection Buffer (LDB)	8.7.5.7	M:8	M:8	
MC11	Support LDB Flush Timer	8.7.5.7	MC10:M	M:8	
MC12	Support of TARP Echo Function	8.7.7	O	i	
MC13	Support for LDB Entry Timer	8.7.5.7	M:8	M:8	

O.1 = One or more of these items must be supported.

Notes:

4. if an NE has multiple NETs
5. if an ES NE
6. if a Level 1 IS NE
7. if a Level 2 IS NE
8. if an IS NE

C.8.5 TARP Processor Management

Index	Functionality/Description	Reference	Status	Profile	Support
PM1	Capable of selectively disabling the TARP propagation by link/adjacency	8.7.6	MC9:M	M	
PM2	Capable of provisioning tar-lif	8.7.6	M	M	
PM3	Capable of provisioning tar-pro	8.7.6	M	M	
PM4	Not capable of provisioning remaining TARP PDU fields	8.7.6	M	M	
PM5	Capable of provisioning TARP timers	8.7.6	M	M	
PM6	Capable of displaying TDC via the local WS at a minimum	8.7.6	MC3:M	MC3:M	
PM7	Capable of displaying LDB via the local WS at a minimum	8.7.6	MC10:M	MC10:M	
PM8	Capable of displaying TARP sequence number via the local WS at a minimum	8.7.6	M	M	
PM9	Capable of manually flushing TDC	8.7.6	MC3:M	MC3:M	
PM10	Capable of manually flushing LDB	8.7.6	MC10:M	MC10:M	
PM11	Capable of manually provisioning TDC	8.7.6	MC3:M	MC3:M	
PM12	Capable of manually provisioning LDB	8.7.6	MC10:M	MC10:M	
PM13	Capable of manually provisioning TARP sequence number	8.7.6	M	M	
PM14	Capable of disabling of all TARP functions	8.7.6	M	M	
PM15	Capable of disabling of TARP propagation function	8.7.6	MC9:M	MC9:M	
PM16	Capable of disabling of TARP origination functions	8.7.6	M	M	
PM17	Capable of disabling of the TDC	8.7.6	MC3:M	MC3:M	
PM18	Capable of manually generating TARP requests	8.7.6	M	M	
PM19	Capable of manually provisioning a TARP adjacency	8.7.8	O	i	
PM20	Capable of manually provisioning the LDB entry timer	8.7.5.7	M	M	
PM21	Capable of manually provisioning the LDB flush timer	8.7.5.7	M	M	

C.8.5.1 LDB Entry Timer Parameters

Index	Functionality/Description	Reference	Status	Profile	Support
LET1	The LDB entry timer shall be settable within a range of 1 to 10 minutes	8.7.5.7	M	M	
LET2	The default value of the LDB entry timer shall be 5 minutes	8.7.5.7	M	M	

C.8.5.2 LDB Flush Timer Parameters

Index	Functionality/Description	Reference	Status	Profile	Support
LFT1	The LDB flush timer shall be settable within a range of 0 to 1440 minutes	8.7.5.7	M	M	
LFT2	The default value of the LDB flush timer shall be 5 minutes	8.7.5.7	M	M	

C.8.5.3 Provisionable TARP PDU Fields

Index	PDU Field	Reference	Status	Profile/ Supported Range & default	Support
PTF1	tar-lif	8.7.2.1, 8.7.6	M	M/ 0 to 65,535 hops & default of 100 hops	
PTF2	tar-pro	8.7.2.3, 8.7.6	M	M/ 00' to FF' (hex) & default of FE' (hex)	

Appendix D: SONET Operations Communications Upper Layers Protocol Profile

D.1 Introduction

This appendix provides an upper layers profile for Session Layer, Presentation Layer, and ACSE protocols to be used by SONET NEs across the SONET OS/NE and NE/NE interfaces. This profile is intended to be used together with the SONET Lower Layers Profile (see Appendix C) and with future SONET Application Service Element (ASE)-specific profiles. The specific applications that have been considered when defining the SONET Upper Layers Profile are Interactive Class (CMISE and TL1 ASEs) and File-oriented Class (FTAM ASE).¹ While most of the SONET Upper Layers Profile is common across the above applications, there are some ASE-specific differences that have been noted.

D.2 Source Documents

When defining the SONET Upper Layers Profile, the source documents listed in the subsections were used.

D.2.1 Base Standards

- ISO/IEC 8650:1988: *Information technology - Open Systems Interconnection - Protocol specification for the Association Control Service Element*
- ISO/IEC 8823-1:1988: *Information processing systems - Open Systems Interconnection - Connection oriented presentation protocol specification - Part 1: Protocol specification*
- ISO/IEC 8327-1:1988: *Information processing systems - Open Systems Interconnection - Connection oriented session protocol specification - Part 1: Protocol specification*
- ISO/IEC 10040:1992: *Information technology - Open Systems Interconnection - Systems Management Overview*

D.2.2 PICS Proforma

- ISO/IEC DIS 8650-2: *Information technology - Open Systems Interconnection - Protocol*

1. If another application layer protocol is chosen for the Remote Login application (e.g., Virtual Terminal), then the SONET Upper Layers protocol profile may need to be updated accordingly.

specification for the Association Control Service Element - Part 2: Protocol Implementation Conformance Statement (PICS) proforma

ISO/IEC DIS 8823-2: *Information technology - Open Systems Interconnection - Connection oriented presentation protocol specification - Part 2: Protocol Implementation Conformance Statement (PICS) proforma*

ISO/IEC DIS 8327-2: *Information technology - Open Systems Interconnection - Basic connection oriented session protocol specification - Part 2: Protocol Implementation Conformance Statement (PICS) proforma*

D.2.3 International Standardized Profiles

ISO/IEC ISP 10607-1: *Information technology - International Standardized Profiles AFTnn - File Transfer, Access and Management - Part 1: Specification of ACSE, Presentation and Session Protocols for use by FTAM*

ISO/IEC ISP 11183-1: *Information technology - International Standardized Profiles AOMIn OSI Management- Management Communications - Part 1: Specification of ACSE, Presentation and Session Protocols for the use by ROSE and CMISE*

ISO/IEC DISP 11188-1: *Information technology - International Standardized Profile - Common Upper Layer Requirements - Part 1: Basic connection oriented requirements*

D.2.4 Bellcore Requirements

Section 8 of this GR

GR-828-CORE

GR-1250

D.3 Goals of SONET Upper Layers Profile

The goals of the SONET Upper Layers Profile are as follows:

- Define an upper layers profile that is common to all SONET applications (whenever practical).
- Remain consistent with ISPs (when possible without violating the above goal).

These goals are consistent with the direction of ISO standards work as reflected in DISP 11188-1.

D.4 Structure of SONET Upper Layers Profile

The SONET Upper Layers Profile is contained in Sections D.6 (ACSE), D.7 (Presentation Layer), and D.8 (Session Layer). Within each of these sections, the following material (specific to the given protocol layer) is contained:

- a subsection called "*Additions Beyond Existing ISP Requirements*" that points out the items for which conformance to ISPs 10607-1 or 11183-1 is not by itself sufficient to meet the SONET Upper Layers Profile,
- a subsection with tables taken from the ISO PICS proforma (with their original table numbers as designated in the PICS proforma),
- a "profile" column added to the above tables indicating the constraints that are imposed by the SONET Upper Layers Profile.

Note that some conditional parts of the existing ISPs (whose conditions do not apply to SONET operations as defined by the Bellcore requirements documents referenced in Section D.2) have not been included in the SONET Upper Layers Profile. For example, not all of the conditional FTAM abstract syntaxes given by ISP 10607-1 apply to SONET operations.

D.5 Notations Used in the SONET Upper Layers Profile

The notations used by the SONET Upper Layers Profile are essentially the same as those that are used in the ISO PICS Proforma documents. The following is taken directly from those documents (except for the "profile column" which is not present in the PICS Proforma).

D.5.1 Abbreviations

Sts	status column
Spt	support column
Sdr	sender
Rcv	receiver
Pfl	profile column

D.5.2 Status Column

This column indicates the level of support required for conformance to the ISO base standard (for the given protocol layer). The values are as follows:

'm' - mandatory support is required

- 'o' - optional support is permitted for conformance to the base standard. If implemented it must conform to the specifications and restrictions contained in the base standard. These restrictions may affect the optionality of other items
- 'n/a' - the item is not applicable
- 'cn' - the item is conditional (where *n* is the number which identifies the condition which is applicable)
- 'o.n' - the item is optional, but the optionality is qualified (where *n* is the number which identifies the qualification which is applicable).

D.5.3 Profile Column

This column indicates the level of support required for conformance to the SONET Upper Layers Profile (for the given protocol layer). The values are as follows:

- 'm' - mandatory support of the feature is required, however, it is not a requirement that the feature be used in all instances of communication, unless mandated by the base standard or stated otherwise in this profile
- 'm(n)' - mandatory support is required as clarified by Note *n*
- 'c(n)' - the item is conditional (where *n* is the Note number which identifies the condition which is applicable)
- 'i' - the item is out of scope for the profile, however, implementation of that item is not precluded; for the purposes of the this profile, the implementation or non-implementation of that item is ignored
- 'x' - the item is excluded (i.e., prohibited) for the profile.

D.5.4 Support Column

The 'Support' column can be completed by the supplier or implementor to indicate the level of implementation of each feature. The proforma has been designed such that the only entries required in the 'Support' column are:

- 'Y' - yes, the feature has been implemented
- 'N' - no, the feature has not been implemented
- '—' - not applicable.

D.5.5 PICS Numbers

Each line within the PICS proforma which requires implementation detail to be entered is numbered at the left-hand edge of the line. This numbering is included as a means of uniquely identifying all possible implementation details within the PICS proforma.

The means of referencing individual responses should be to specify the following sequence:

- a. a reference to the smallest subclause enclosing the relevant item
- b. a solidus character, '/'
- c. the reference number of the row in which the response appears
- d. if, and only if, more than one response occurs in the row identified by the reference number, then each possible entry is implicitly labeled a, b, c, etc., from left to right, and this letter is appended to the sequence.

D.6 SONET Upper Layers Profile: ACSE

D.6.1 Additions Beyond Existing ISP Requirements

The numbers used below refer to the PICS Proforma numbers used in Section D.6.2.

A.5.1 Association establishment

This profile requires that both the Initiator and Responder roles always be supported. ISP 10607-1 and ISP 11183-1 do not always require support of both roles.

A.9 Supported APDU parameters

Application Entity Title

When used with the CMISE or FTAM ASEs, this profile requires support for sending the following parameters in the AARQ APDU: Calling AP title, Calling AE qualifier, Called AP title, and Called AE qualifier. Similarly, it requires (for the FTAM ASE only) sending the "Responding AP title" and the "Responding AE qualifier" parameters in the AARE APDU.

For ISP 10607-1, these parameters are also mandatory.

For ISP 11183-1, these parameters are optional.

A.10.1 AE title name form

When used with the CMISE or FTAM ASEs, this profile requires support for sending Form 1 AE title name forms. For the FTAM ASE, it also requires support for sending Form 2.

For ISP 10607-1, Form 2 is required and Form 1 is optional.

For ISP 11183-1, both forms are optional, i.e., use of AE title is optional.

D.6.2 Profile Tables

Note that the tables below use the PICS numbers of the ISO PICS Proforma. Only those tables from the PICS Proforma that apply to the SONET Upper Layers Profile are provided.

A.5 Supported roles

A.5.1 Association establishment

	Role	Sts	Pfl	Spt	Associated predicate
1	Initiator	o.01	m		A-CON/initiator
2	Responder	o.01	m		A-CON/responder

o.01: a conforming implementation shall support at least one of the roles.

A.5.2 Normal release

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.02	m		A-REL/requester
2	Acceptor	o.02	m		A-REL/acceptor

o.02: a conforming implementation shall support at least one of the roles.

A.6 Protocol mechanisms

	Protocol mechanism	Sts	Pfl	Spt	Associated predicate
1	Normal mode	o.03	m		
2	X.410-1984 mode	o.03	i		
3	Rules for extensibility	m	m		
4	Support operation of Session version 2	o	m		S-O-SESS-V2

o.03: either Normal mode or X.410-1984 mode or both shall be supported. If only X.410-1984 mode is supported, then the remainder of the proforma shall be ignored.

A.7 Functional units

	ACSE functional units	Sts	Pfl	Spt	Associated predicate
1	Kernel	m	m		
2	Authentication	o	i		A-FU(AU)

A.8 Supported APDUs

	APDU	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	A-associate-request (AARQ)	c01	m		c02	m	
2	A-associate-response (AARE)	c02	m		c01	m	
3	A-release-request (RLRQ)	c03	m		c04	m	
4	A-release-response (RLRE)	c04	m		c03	m	
5	A-abort (ABRT)	c05	m		c05	m	

c01: if [A-CON/initiator] then m else n/a
c02: if [A-CON/responder] then m else n/a
c03: if [A-REL/initiator] then m else n/a
c04: if [A-REL/responder] then m else n/a
c05: if [S-O-SESS-V2] then m else n/a

A.9 Supported APDU parameters

Note: Applications may place further constraints on the use of APDU "User information" parameters.

A.9.1 A-associate-request (AARQ)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol version	c06	m(2)		c02	m(2)	
2	Application context name	c01	m(6)		c02	m	
3	Calling AP title	c06	c(5)		c02	m(1)	
4	Calling AE qualifier	c06	c(5)		c02	m(1)	
5	Calling AP-invocation-identifier	c06	i		c02	m	
6	Calling AE-invocation-identifier	c06	i		c02	m	
7	Called AP title	c06	c(5)		c02	m(1)	
8	Called AE qualifier	c06	c(5)		c02	m(1)	
9	Called AP-invocation-identifier	c06	i		c02	m	
10	Called AE-invocation-identifier	c06	i		c02	m	
11	ACSE-requirements	c07	i		c08	m(3)	
12	Authentication-mechanism-name	c07	i		c08	m(3)	
13	Authentication-value	c07	i		c08	m(3)	
14	Implementation information	c06	i		c02	m	
15	User information	c06	c(5)		c02	m	

c01: if [A-CON/initiator] then m else n/a

- c02: if [A-CON/responder] then m else n/a
c06: if [A-CON/initiator] then o else n/a
c07: if [A-CON/initiator and A-FU(AU)] then m else n/a
c08: if [A-CON/responder and A-FU(AU)] then m else n/a
(1) Both forms shall be static mandatory/dynamically optional for receiving (see Table A.10.1).
(2) The default value "version 1" is defined in the abstract syntax definition of ACSE APDUs in ISO/IES 8650. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.
(3) If the authentication FU is not supported, based on the extensibility rules, these tagged values shall be received and ignored. The "Authentication-mechanism-name" and "Authentication-value" fields shall only be present if the "ACSE-requirements" field includes the authentication FU. The "Authentication-mechanism-name" field shall be present if "Authentication-value" is of type ANY DEFINED BY.
(5) if CMISE or FTAM ASE then m else i
(6) Values for the application context name parameter are specified in the application layer profiles.

A.9.2 A-associate-response (AARE)

Parameter	Sending			Receiving		
	Sts	Pfl	Spt	Sts	Pfl	Spt
1 Protocol version	c09	m(2)		c01	m(2)	
2 Application context name	c02	m(6)		c01	m	
3 Responding AP title	c09	c(7)		c01	m(1)	
4 Responding AE qualifier	c09	c(7)		c01	m(1)	
5 Responding AP-invocation-identifier	c09	i		c01	m	
6 Responding AE-invocation-identifier	c09	i		c01	m	
7 Result	c02	m		c01	m	
8 Result source diagnostic	c10	m		c11	m	
9 ACSE-requirements	c08	i		c07	m(3)	
10 Authentication-mechanism-name	c08	i		c07	m(3)	
11 Authentication-value	c08	i		c07	m(3)	
12 Implementation information	c09	i		c01	m	
13 User information	c09	c(5)		c01	m	

Note: notes from previous table (A.9.1) also apply to this table (A.9.2).

- c09: if [A-CON/responder] then o else n/a
c10: if [A-CON/responder] then (if [A-FU(AU)] then m (with a value range of 11 to 14) else o (with a value range of 1 to 10)) else n/a
c11: if [A-CON/initiator] then (if [A-FU(AU)] then m (with a value range of 11 to 14) else o (with a value range of 1 to 10)) else n/a
(7) if FTAM ASE then m else i

A.9.3 A-release-request (RLRQ)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Reason	c12	i		c04	m	
2	User information	c12	c(1)		c04	m	

c04: if [A-REL/acceptor] then m else n/a

c12: if [A-REL/requester] then o else n/a

(1) if FTAM ASE then m else i

A.9.4 A-release-response (RLRE)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Reason	c13	i		c03	m	
2	User information	c13	c(1)		c03	m	

c03: if [A-REL/initiator] then m else n/a

c13: if [A-REL/acceptor] then o else n/a

(1) if FTAM ASE then m else i

A.9.5 A-abort (ABRT)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Abort source	m	m		m	m	
2	Diagnostic	c6	i		c6	i	
3	User information	o	m		m	m	

c6: if [A-FU(AU)] then m else n/a

A.10 Supported parameter forms

A.10.1 AE title name form

	Syntax form	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Form 1 (Directory name)	o.04	m		m	m(1)	
2	Form 2 (Object identifier and integer)	o.04	c(2)		m	m(1)	

o.04: a conforming implementation shall support at least one of the forms.

(1) Both forms should be static mandatory/dynamically optional for receiving.

(2) if FTAM ASE then m else i

D.7 SONET Upper Layers Profile: Presentation Layer

D.7.1 Additions Beyond Existing ISP Requirements

The numbers used below refer to the PICS Proforma numbers used in Section D.7.2.

A.7.3 CPR PPDU

For this profile, implementations shall support sending the "Provider reason" parameter in the CPR PPDU.

For ISP 11183-1, the same requirement applies.

For ISP 10607-1, support of the "Provider reason" parameter is optional.

A.7.5 ARP PPDU

For this profile, implementations shall support sending the "Abort reason" parameter in the ARP PPDU. The "Event identifier" parameter shall be present if the "Abort reason" parameter is set with the value 2, 3, 4, 5 or 6.

For ISP 11183-1, the same requirement applies.

For ISP 10607-1, support of the "Abort reason" parameter in the ARP PPDU is optional and support of the "Event identifier" parameter is out of scope.

D.7.2 Profile Tables

Note that the tables below use the PICS numbers of the ISO PICS Proforma. Only those tables from the PICS Proforma that apply to the SONET Upper Layers Profile are provided.

A.5 Protocol mechanisms and functional units

A.5.1 Protocol mechanisms

	Protocol mechanism	Sts	Pfl	Spt	Associated predicate
1	X.410-1984 mode	o.01	i		P-MODE(X.410)
2	Normal	o.01	m		P-MODE(NORMAL)

o.01: either Normal mode or X.410 (1984) mode or both shall be supported.

A.5.2 Functional units

	Presentation functional units	Sts	Pfl	Spt	Associated predicate
1	Kernel	m	m		
2	Presentation Context Management	c00	i		P-FU(CM)
3	Presentation Context Restoration	c01	i		P-FU(CR)
	Pass through to Session functional units				
4	Negotiated Release	o	Note 1		S-FU(NR)
5	Half Duplex	o.02	Note 1		S-FU(HD)
6	Duplex	o.02	Note 1		S-FU(FD)
7	Expedited Data	o	Note 1		S-FU(EX)
8	Typed Data	o	Note 1		S-FU(TD)
9	Capability Data Exchange	c02	Note 1		S-FU(CD)
10	Minor Synchronize	o	Note 1		S-FU(SY)
11	Symmetric Synchronize	o	Note 1		S-FU(SS)
12	Major Synchronize	o	Note 1		S-FU(MA)
13	Resynchronize	o	Note 1		S-FU(RESYN)
14	Exceptions	c03	Note 1		S-FU(EXCEP)
15	Activity Management	o	Note 1		S-FU(ACT)

Note 1: See Section D.8 (Session Layer Profile).

o.02: pass through for at least one of the Session functional units Duplex and Half Duplex shall be supported.

c00: if [P-MODE(NORMAL)] then o else n/a

c01: if [P-FU(CM)] then o else n/a

c02: if [S-FU(ACT)] then o else n/a

c03: if [S-FU(HD)] then o else n/a

A.6 Elements of procedure related to the PICS

A.6.1 Kernel functional unit

A.6.1.1 Supported roles

A.6.1.1.1 Presentation connection

	Role	Sts	Pfl	Spt	Associated predicate
1	Initiator	o.03	m		P-CON/initiator
2	Responder	o.03	m		P-CON/responder

o.03: a conforming implementation shall support at least one of the roles.

A.6.1.1.2 Normal data

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.04	m		P-DATA/requester
2	Acceptor	o.04	m		P-DATA/acceptor

o.04: a conforming implementation shall support at least one of the roles.

A.6.1.1.3 Orderly release

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.05	m		P-REL/requester
2	Acceptor	o.05	m		P-REL/acceptor

o.05: a conforming implementation shall support at least one of the roles.

A.6.1.2 Supported PPDU's associated with the kernel services

	PPDU	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Connect presentation (CP)	c04	m(1)		c05	m(1)	
2	Connect presentation Accept(CPA)	c05	m		c04	m	
3	Connect presentation Reject (CPR)	c05	m		c04	m	
4	Abnormal release provider (ARP)	m	m		m	m	
5	Abnormal release user (ARU)	o	m		m	m	
6	Presentation data (TD)	c06	m		c07	m	

c04: if [P-CON/initiator] then m else n/a

c05: if [P-CON/responder] then m else n/a

c06: if [P-DATA/requester] then m else n/a

c07: if [A-DATA/acceptor] then m else n/a

(1) includes Cptype (see ISP 11183-1, Item B.3.1/2)

A.7 Supported PPDU parameters

Note: Applications may place constraints on the encoding of PPDU "User data parameters", e.g., see Section 8.3.7.5 for constraints specified for TL1.

A.7.1 Connect presentation (CP) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling presentation selector	c10	i		c05	m	
2	Called presentation selector	c10	m		c05	m	
3	Mode selector	c04	m		c05	m	
4	Presentation context definition list	c10	m(1)		c05	m(1)	
5	Default context name	c10	i		c05	m	
6	Protocol version	c04	m(2)		c05	m(2)	
7	Presentation requirements	c10	i		c05	m	
8	User session requirements	c11	i		c05	m	
9	User data	c10	m		c05	m	

c04: if [P-CON/initiator] then m else n/a

c05: if [P-CON/responder] then m else n/a

c10: if [P-CON/initiator] then o else n/a

c11: if [P-CON/initiator and P-FU(CM)] then o else n/a

(1) A conforming implementation shall encode presentation context identifiers in the range 0 to 32,767. For selection of odd or even value, see ISO 8823, 6.2.2.7 and 6.5.2.1.

(2) The default value "version 1" is defined in the structure of SS-user data definition in clause 8.2 of ISO/IEC 8823. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

A.7.2 Connect presentation accept (CPA) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Responding presentation selector	c12	i		c04	m	
2	Mode selector	c05	m		c04	m	
3	Presentation context definition result list	c05	m(1)		c13	m(1)	
4	Protocol version	c05	m(2)		c04	m(2)	
5	Presentation requirements	c12	i		c14	m	
6	User session requirements	c12	i		c15	m	
7	User data	c12	m		c04	m	

c04: if [P-CON/initiator] then m else n/a

c05: if [P-CON/responder] then m else n/a

- c12: if [P-CON/responder and P-FU(CM)] then o else n/a
c13: if [P-CON/initiator and A.7.1/4a] then m else n/a
c14: if [P-CON/initiator and A.7.1/7a] then m else n/a
c15: if [P-CON/initiator and A.7.1/8a] then m else n/a
(1) A conforming implementation shall encode presentation context identifiers in the range 0 to 32,767.
(2) The default value "version 1" is defined in the structure of SS-user data definition in clause 8.2 of ISO/IEC 8823. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

A.7.3 Connect presentation reject (CPR) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Responding presentation selector	c12	i		c04	m	
2	Presentation context definition result list	c12	m(1)		c13	m(1)	
3	Protocol version	c12	m(3)		c04	m(3)	
4	Default context result	c12	i		c16	m	
5	Provider reason	c12	m		c04	m	
6	User data	c12	m(2)		c04	m	

- c04: if [P-CON/initiator] then m else n/a
c12: if [P-CON/responder and P-FU(CM)] then o else n/a
c16: if [P-CON/initiator and A.7.1/5a] then m else n/a
(1) The "Presentation context definition result list" is required if the "Provider reason" parameter is absent. If the "Provider reason" is present, then the "Presentation context definition result list" is optional.
(2) Is not present if the connection is rejected by the Presentation service provider.
(3) The default value "version 1" is defined in the structure of SS-user data definition in clause 8.2 of ISO/IEC 8823. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

A.7.4 Abnormal release user (ARU) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Presentation context identifier list	c17	m	m	m	m	m
2	User data	c18	m	m	m	m	m

- c17: if [A.6.1.2/5a] then (if [P-FU(CM) and A.7.5/2a or A.7.2/4a or P-CON/responder] then m else o) else n/a
c18: if [A.6.1.2/5a] then o else n/a

A.7.5 Abnormal release provider (ARP) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Abort reason	m	m		m	m	
2	Event identifier	o	m(1)		m	m(1)	

(1) Mandatory if "Provider reason" parameter is set with the following values: 2, 3, 4, 5, 6.

A.7.6 Presentation data (TD) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	User data	c06	m		c07	m	

c06: if [P-DATA/requester] then m else n/a

c07: if [A-DATA/acceptor] then m else n/a

A.8 Support of syntaxes

A.8.1 Transfer syntaxes supported

	Type	Detail	Support		Reference to definition	Reference to restriction
			Pfl	Spt		
1	Object identifier	{joint-iso-ccitt asn1(1) basic-encoding(1)}	c(1)		ISO/IEC 8825	see note 3
2	Object identifier	{1 3 17 104 12 2 bellcoreSONETSyntax(1)}	c(2)		Section 8.3.7.5	

(1) if CMISE or FTAM ASE then m else n/a

(2) if TL1 ASE then m else n/a

(3) For further restrictions on the use of ASN.1 Basic Encoding, see ISO/IEC DISP 11188-1, Common Upper Layer Requirements.

A.8.2 Abstract syntaxes supported

	Type	Detail	Support		Reference to definition	Reference to restriction
			Pfl	Spt		
1	Object identifier	{joint-iso-ccitt association-control(2) abstract-syntax(1) apdus(0) version1(1)}	m		ISO 8650	
2	Object identifier	{joint-iso-ccitt ms(9) cmip(1) cmip-pci(1) abstractSyntax(4)}	c(1)		ISO 9596-1	
3	Object identifier	{joint-iso-ccitt ms(9) smo(0) negotiationAbstractSyntax(1) version1(1)}	c(1)		ISO 10040	
4	FTAM PCI	{iso standard 8571 abstract-syntax(2) ftam-pci(1)}	c(2)		ISP 10607-1	
5	FTAM unstructured text	{iso standard 8571 abstract-syntax(2) unstructured-text(3)}	c(2)		ISP 10607-1	
6	FTAM unstructured binary	{iso standard 8571 abstract-syntax(2) unstructured-binary(4)}	c(2)		ISP 10607-1	
7	Object identifier	{1 13 17 104 11 2 bellcoreSONETSyntax(1)}	c(3)		Section 8.3.7.5	

Note: ISP 10607-1 conditionally requires seven additional abstract syntaxes for the FTAM ASE (not listed here), however, their conditions are false for the SONET File Transfer application.

- (1) if CMISE ASE then m else n/a
- (2) if FTAM ASE then m else n/a
- (3) if TL1 ASE then m else n/a

D.8 SONET Upper Layers Profile: Session Layer

D.8.1 Additions Beyond Existing ISP Requirements

The numbers used below refer to the PICS Proforma numbers used in Section D.8.2.

A.8.9 Abort (AB) SPDU

Reflect Parameter Values

For this profile, the "Reflect Parameter Values" parameter in the AB SPDU is conditionally present only if the "Transport Disconnect" value is "protocol error".

For ISP 11183-1, the same requirement applies.

For ISP 10607-1, the presence of "Reflect Parameter Values" parameter in the AB SPDU is optional.

D.8.2 Profile Tables

Note that the tables below use the PICS numbers of the ISO PICS Proforma. Only those tables from the PICS Proforma that apply to the SONET Upper Layers Profile are provided.

A.3 ISO 8327 protocol versions implemented

	Version	Sts	Pfl	Spt	Associated predicate
1	Version 1	o.1	i		S-V1
2	Version 2	o.1	m		

o.1: At least one version shall be implemented.

A.6 Supported functional units and protocol mechanisms

A.6.1 Functional units

	Functional Unit	Sts	Pfl	Spt	Associated predicate
1	Kernel	m	m		
2	Negotiated Release	o	i		S-FUN(NR)
3	Half Duplex	o.2	i		S-FU(HD)
4	Duplex	o.2	m		S-FU(FD)
5	Expedited Data	o	i		S-FU(EX)
6	Typed data	o	i		S-FU(TD)
7	Capability Data Exchange	c1	i		S-FU(CD)
8	Minor Synchronize	o	i		S-FU(SY)
9	Symmetric Synchronize	o	i		S-FU(SS)
10	Major Synchronize	o	i		S-FU(MA)
11	Resynchronize	o	i		S-FU(RESYN)
12	Exceptions	c2	i		S-FU(EXCEP)
13	Activity management	o	i		S-FU(ACT)

o.2: At least one of the functional units Duplex and Half Duplex shall be implemented.

c1: if [S-FU(ACT)] then o else n/a

c2: if [S-FU(HD)] then o else n/a

A.6.2 Protocol mechanisms

	Mechanism	Sts	Pfl	Spt	Associated predicate
1	Use of transport expedited data (Extended control Quality Of Service)	o	i		S-EXP/T
2	Reuse of transport connection	o	i		S-REUSE/T
3	Basic concatenation	m	m		
4	Extended concatenation (sending)	o	i		
5	Extended concatenation (receiving)	o	i		S-XCONC/RCV
6	Segmenting (sending)	o	i		S-SEG/SDR
7	Segmenting (receiving)	o	i		S-SEG/RCV
8	Max size of SS-user-data < or = 512	o	x		S-MAXSIZE/512
9	Max size of SS-user-data < or = 10240	o	see note 1		S-MAXSIZE/10240
10	Max size of SS-user-data < or = 9	o	x		S-MAXSIZE/9

- (1) at least 10240 octets is the required Max size of SS-user-data. It is a Bellcore objective to be able to support a Max size of 65,535 octets (see Section 8.3.5).

A.7 Elements of procedure related to the PICS

A.7.1 Kernel functional unit

A.7.1.1 Supported roles for the Kernel functional unit services

A.7.1.1.1 Session Connection

Does the implementation support the Session Connection as:

	Role	Sts	Pfl	Spt	Associated predicate
1	Initiator	o.3	m		S-CON/initiator
2	Responder	o.3	m		S-CON/responder

- o.3: a conforming implementation must support at least one of the above roles.

A.7.1.1.2 Orderly Release

Does the implementation support the Orderly Release as:

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.4	m		S-REL/requester
2	Acceptor	o.4	m		S-REL/acceptor

o.4: a conforming implementation must support at least one of the above roles.

A.7.1.1.3 Normal Data Transfer

Does the implementation support the Normal Data Transfer as:

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.5	m		S-DATA/requester
2	Acceptor	o.5	m		S-DATA/acceptor

o.5: a conforming implementation must support at least one of the above roles.

A.7.1.2 Support for the SPDUs associated with the Kernel services

	SPDU	Sending			Receiving			Associated predicate	
		Sts	Pfl	Spt	Sts	Pfl	Spt	Sending	Receiving
1	Connect (CN)	c3	m		c4	m			
2	Overflow Accept (OA)	c5	i		c6	i		S-OA/SDR	S-OA/RCV
3	Connect Data Overflow (CDO)	c6	i		c5	i		S-CDO/SDR	S-CDO/RCV
4	Accept (AC)	c4	m		c3	m			
5	Refuse (RF)	c4	m		c3	m			
6	Finish (FN)	c7	m		c8	m			
7	Disconnect (DN)	c8	m		c7	m			
8	Abort (AB)	m	m		m	m			
9	Abort Accept (AA)	o	i		o	i			
10	Data Transfer (DT)	c9	m		c10	m			
11	Prepare (PR)	c11	i		c11	i		S-PR/SDR	S-PR/RCV

c3: if [S-CON/initiator] then m else n/a
c4: if [S-CON/responder] then m else n/a
c5: if [S-V1 or (NOT S-CON/responder)] then n/a else if [NOT S-MAXSIZE/10240] then m else o
c6: if [S-V1 or (NOT S-CON/initiator)] then n/a else if [NOT S-MAXSIZE/10240] then m else o
c7: if [S-REL/requester] then m else n/a
c8: if [S-REL/acceptor] then m else n/a
c9: if [S-DATA/requester] then m else n/a
c10: if [S-DATA/acceptor] then m else n/a
c11: if [S-V1] then n/a else if [NOT S-MAXSIZE/9 and S-EXP/T] then m else o

A.7.1.3 Support for the SPDUs associated with Token Exchange

	SPDU	Sending			Receiving			Comment
		Sts	Pfl	Spt	Sts	Pfl	Spt	
1	Give Token (GT)	m	m		m	m		
2	Please Token (PT)	m	i (1)		m	m		

(1) According to the base standard basic concatenation rules, the conditions under which PT would be used are all out-of-scope for this profile.

Note: Tables A.7.2 through A.7.3 of the PICS Proforma are not applicable to this profile.

A.7.4 Duplex functional unit [Associated predicate: S-FU(FD)]

No additional SPDUs (this clause is present for completeness).

Note: Tables A.7.5 through A.7.13 of the PICS Proforma are not applicable to this profile

A.8 Supported SPDU-parameters

A.8.1 Connect (CN) SPDU

	PGI "Connection Identifier"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling SS-user Reference	c46	i		c4	m	
2	Common Reference	c46	i		c4	m	
3	Additional Reference Information	c46	i		c4	m	

c4: if [S-CON/responder] then m else n/a

c46: if [S-CON/initiator] then o else n/a

A.8.1.2 Connect/Accept Item

A.8.1.2.1 Connect/Accept Item parameters

Important Remark: If presence of the PGI "Connect/accept Item" is supported (see clause A.8.1.2.2) then presence of Protocol Options and Version Number parameters must be supported.

	PGI "Connection/Accept Item"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol Options	c47	m		c48	m	
2	TSDU maximum size	c49	i		c50	m	
3	Version Number	c51	m		c52	m	
4	Initial Serial Number	c53	i		c54	m	
5	Token Setting Item	c46	i		c55	m	
6	Second Initial Serial Number	c56	i		c57	i	

c46: if [S-CON/initiator] then o else n/a

c47: if [NOT S-CON/initiator] then n/a else if [S-XCONC/RCV] then m else o

c48: if [NOT S-CON/responder] then n/a else if [S-XCONC/RCV] then m else o

c49: if [NOT S-CON/initiator] then n/a else if [S-SEG/SDR or S-SEG/RCV] then m else o

c50: if [NOT S-CON/responder] then n/a else if [S-SEG/SDR or S-SEG/RCV] then m else o

c51: if [NOT S-CON/initiator] then n/a else if [NOT S-V1] then m else o

c52: if [NOT S-CON/responder] then n/a else if [NOT S-V1] then m else o

c53: if [NOT S-CON/initiator] then n/a
else if [(S-FU(SY) or S-FU(MA) or S-FU(SS) or S-FU(RESYN)) and NOT S-FU(ACT)] then m
else o

c54: if [NOT S-CON/responder] then n/a
else if [(S-FU(SY) or S-FU(MA) or S-FU(SS) or S-FU(RESYN)) and NOT S-FU(ACT)] then m
else o

c55: if [S-CON/responder] then o else n/a

c56: if [NOT S-CON/initiator] then n/a else if [S-FU(SS) and NOT S-FU(ACT)] then m else o

c57: if [NOT S-CON/responder] then n/a else if [S-FU(SS) and NOT S-FU(ACT)] then m else o

A.8.1.2.2 Presence of Connect/Accept Item

	Presence of "Connection/Accept Item"	Sts	Pfl	Spt
1	Sending	c58	m	
2	Receiving	c59	m	

c58: if [NOT S-CON/initiator] then n/a
else if [A8.1.2.1/1a or A.8.1.2.1/2a or A.8.1.2.1/3a or A.8.1.2.1/4a or A.8.1.2.1/5a or A.8.1.2.1/6a]
then m else o

c59: if [NOT S-CON/responder] then n/a
else if [A8.1.2.1/1b or A.8.1.2.1/2b or A.8.1.2.1/3b or A.8.1.2.1/4b or A.8.1.2.1/5b or A.8.1.2.1/6b]
then m else o

A.8.1.3 Single Items

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Session User Requirements	c60	m		c61	m	
2	Calling Session Selector	c46	i		c4	m	
3	Called Session Selector	c3	m(1)		c4	m	
4	Data Overflow	c6	i		c5	i	
5	User Data	c3	m		c4	m	
6	Extended User Data	c62	m		c63	m	

- c3: if [S-CON/initiator] then m else n/a
c4: if [S-CON/responder] then m else n/a
c5: if [S-V1 or (NOT S-CON/responder)] then n/a
else if [NOT S-MAXSIZE/10240] then m else o
c6: if [S-V1 or (NOT S-CON/initiator)] then n/a
else if [NOT S-MAXSIZE/10240] then m else o
c46: if [S-CON/initiator] then o else n/a
c60: if [NOT S-CON/initiator] then n/a
else if [S-FU(HD) and S-FU(SY) and S-FU(ACT) and S-FU(CD) and S-FU(EXCEP)] then o
else m
c61: if [NOT S-CON/responder] then n/a
else if [S-FU(HD) and S-FU(SY) and S-FU(ACT) and S-FU(CD) and S-FU(EXCEP)] then o
else m
c62: if [S-V1 or (NOT S-CON/initiator)] then n/a
else if [NOT S-MAXSIZE/512] then m else o
c63: if [S-V1 or (NOT S-CON/responder)] then n/a
else if [NOT S-MAXSIZE/512] then m else o
(1) The value of the Called Session Selector parameter shall be ASCII "SS" (i.e., "5353" hex) to indicate that the ISO Presentation Layer is being run over the ISO Session Layer.

Note: Tables A.8.2 through A.8.3 of the PICS Proforma are not applicable to this profile.

A.8.4 Accept (AC) SPDU

A.8.4.1 Connect Identifier

	PGI "Connection Identifier"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling SS-user Reference	c55	i		c3	m	
2	Common Reference	c55	i		c3	m	
3	Additional Reference Information	c55	i		c3	m	

- c3: if [S-CON/initiator] then m else n/a
c55: if [S-CON/responder] then o else n/a

A.8.4.2 Connect/Accept Item

A.8.4.2.1 Connect/Accept Item parameters

Important Remark: If presence of the PGI "Connect/accept Item" is supported (see clause A.8.4.2.2) then presence of Protocol Options and Version Number parameters must be supported.

	PGI "Connection/Accept Item"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol Options	c48	m		c47	m	
2	TSDU maximum size	c50	i		c49	m	
3	Version Number	c52	m		c51	m	
4	Initial Serial Number	c54	i		c53	m	
5	Token Setting Item	c55	i		c46	m	
6	Second Initial Serial Number	c57	i		c56	i	

A.8.4.2.2 Presence of Connect/Accept Item

	Presence of "Connection/Accept Item"	Sts	Pfl	Spt
1	Sending	c70	m	
2	Receiving	c71	m	

c70: if [NOT S-CON/responder] then n/a
 else if [A8.4.2.1/1a or A.8.4.2.1/2a or A.8.4.2.1/3a or A.8.4.2.1/4a or A.8.4.2.1/5a or A.8.4.2.1/6a]
 then m else o
 c71: if [NOT S-CON/initiator] then m
 else if [A8.4.2.1/1b or A.8.4.2.1/2b or A.8.4.2.1/3b or A.8.4.2.1/4b or A.8.4.2.1/5b or A.8.4.2.1/6b]
 then m else o

A.8.4.3 Single Items

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Token Item	c55	i		c3	m	
2	Session User Requirements	c61	m		c60	m	
3	Enclosure Item	c72	i		c73	i	
4	Calling Session Selector	c4	i		c46	m	
5	Responding Session Selector	c55	i		c3	m	
6	User Data	c4	m		c3	m	

Note: The notes used on Table A.8.1.3 also apply to this table.
 c55: if [S-CON/responder] then o else n/a
 c72: if [S-CON/responder and NOT S-V1] then m else n/a
 c73: if [S-CON/initiator and NOT S-V1] then m else n/a

A.8.5 Refuse (RF) SPDU

A.8.5.1 Connection Identifier

	PGI "Connection Identifier"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling SS-user Reference	c55	i		c3	m	
2	Common Reference	c55	i		c3	m	
3	Additional Reference Information	c55	i		c3	m	

c3: if [S-CON/initiator] then m else n/a
c55: if [S-CON/responder] then o else n/a

A.8.5.2 Single Items

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Transport Disconnect	c74	i		c75	m	
2	Session User Requirements	c(1)	c(1)		c(1)	c(1)	
3	Version Number	c52	m		c51	m	
4	Enclosure Item	c72	i		c73	i	
5	Reason Code	c(2)	c(2)		c(2)	c(2)	

c51:if [NOT S-CON/initiator] then n/a else if [NOT S-V1] then m else o
c52:if [NOT S-CON/responder] then n/a else if [NOT S-V1] then m else o
c72:if [S-CON/responder and NOT S-V1] then m else n/a
c73:if [S-CON/initiator and NOT S-V1] then m else n/a
c74:if [NOT S-CON/responder] then n/a else if [S-REUSE/T] then m else o
c75:if [NOT S-CON/initiator] then n/a else if [S-REUSE/T] then m else o
(1) Shall only be present if the "Reason Code" value = 2.
(2) Mandatory if "Enclosure Item" parameter is present.

A.8.6 Finish (FN) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Transport Disconnect	c76	i		c79	m	
2	Enclosure Item	c77	i		c80	i	
3	User Data	c78	m		c8	m	

c8:if [S-REL/acceptor] then m else n/a
c76:if [NOT S-REL/requester] then n/a else if [S-REUSE/T] then m else o
c77:if [S-REL/requester and NOT S-V1] then m else n/a
c78:if [S-REL/requester] then o else n/a
c79:if [NOT S-REL/acceptor] then n/a else if [S-REUSE/T] then m else o
c80:if [S-REL/acceptor and NOT S-V1] then m else n/a

A.8.7 Disconnect (DN) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Enclosure Item	c80	i		c77	i	
2	User Data	c81	m		c7	m	

c7:if [S-REL/requester] then m else n/a
c77:if [S-REL/requester and NOT S-V1] then m else n/a
c80:if [S-REL/acceptor and NOT S-V1] then m else n/a
c81:if [S-REL/acceptor] then o else n/a

A.8.8 Not Finish (NF) SPDU

Not applicable to this profile.

A.8.9 Abort (AB) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Transport Disconnect	m	m		m	m	
2	Enclosure Item	c83	i		c83	i	
3	Reflect Parameter Values	c(1)	c(1)		o	m	
4	User Data	o	m		m	m	

c83: if [NOT S-V1] then m else n/a
(1) Only sent if Transport disconnect = protocol error.

A.8.10 Abort Accept (AA) SPDU

No Parameter field.

A.8.11 Data Transfer (DT) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Enclosure Item	c84	i		c85	i	
2	User Information Field	c9	m		c10	m	

c9: if [S-DATA/requester] then m else n/a
c10: if [S-DATA/acceptor] then m else n/a
c84: if [S-DATA/requester and S-SEG/SDR] then m else n/a
c85: if [S-DATA/acceptor and S-SEG/RCV] then m else n/a

Note: Tables A.8.12 through A.8.15 of the PICS Proforma are not applicable to this profile

A.8.16 Give Tokens (GT) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Token Item	c98	c(1)		c99	c(1)	
2	Enclosure Item	c83	i		c83	i	
3	User Data	c100	m		c83	m	

c83: if [NOT S-V1] then m else n/a
c98: if [S-FU(NR) or S-FU(HD) or S-FU(SY) or S-FU(MA) or S-FU(ACT)] then o else n/a
c99: if [S-FU(NR) or S-FU(HD) or S-FU(SY) or S-FU(MA) or S-FU(ACT)] then m else n/a
c100: if [NOT S-V1 and A.8.16/1a] then o else n/a
(1) if VT ASE then m else i

Note: Tables A.8.17 through A.8.36 of the PICS Proforma are not applicable to this profile

References

Bellcore Publications

Family of Requirements

- FR-476, *OTGR Section 6: Network Maintenance: Access and Testing* (Bellcore, 1998 Edition). (A subset of OTGR, FR-439.)
- FR-480, *OTGR Section 10: User System Interface* (Bellcore, 1998 Edition). (A subset of OTGR, FR-439.)

Generic Requirements

- GR-20-CORE, *Generic Requirements for Optical Fiber and Optical Fiber Cable*, Issue 2 (Bellcore, July 1998).
- GR-63-CORE, *Network Equipment-Building System (NEBS) Requirements: Physical Protection (a module of LSSGR, FR-64, TSGR, FR-440, and NEBSFR, FR-2063)*, Issue 1 (Bellcore, October 1995).
- GR-78-CORE, *Generic Requirements for the Physical Design and Manufacture of Telecommunications Products and Equipment (a module of RQGR, FR-796 and NEBSFR, FR-2063)*, Issue 1 (Bellcore, September 1997).
- GR-199-CORE, *OTGR Section 12.2: Operations Application Messages - Memory Administration Messages (a module of OTGR, FR-439)*, Issue 2 (Bellcore, November 1996).
- GR-303-CORE, *Integrated Digital Loop Carrier System Generic Requirements, Objectives, and Interface (a module of TSGR, FR-440)*, Issue 2 (Bellcore, December 1998).
- GR-326-CORE, *Generic Requirements for Single-Mode Optical Connectors and Jumper Assemblies*, Issue 2 (Bellcore, December 1996).
- GR-409-CORE, *Generic Requirements for Premises Fiber Optic Cable*, Issue 1 (Bellcore, May 1994).
- GR-418-CORE, *Generic Reliability Assurance Requirements for Fiber Optic Transport Systems (a module of RQGR, FR-796)*, Issue 1 (Bellcore, December 1997).
- GR-436-CORE, *Digital Network Synchronization Plan*, Issue 1 (Bellcore, June 1994) plus Revision 1, June 1996.

- GR-449-CORE, *Generic Requirements and Design Consideration for Fiber Distributing Frames*, Issue 1 (Bellcore, March 1995).
- GR-454-CORE, *Generic Requirements for Supplier Provided Documentation (a module of OTGR, FR-439, LSSGR, FR-64, and TSGR, FR-440)*, Issue 1 (Bellcore, December 1997).
- GR-468-CORE, *Generic Reliability Assurance Requirements for Optoelectronic Devices Used in Telecommunications Equipment (a module of RQGR, FR-796)*, Issue 1 (Bellcore, December 1998).
- GR-472-CORE, *OTGR Section 2.1: Network Element Configuration Management (a module of OTGR, FR-439)*, Issue 2 (Bellcore, November 1996) plus Revision 1, July 1997.
- GR-474-CORE, *OTGR Section 4: Network Maintenance: Alarm and Control for Network Elements (a module of OTGR, FR-439)*, Issue 1 (Bellcore, December 1997).
- GR-487-CORE, *Generic Requirements for Electronic Equipment Cabinets*, Issue 1 (Bellcore, June 1996).
- GR-496-CORE, *SONET Add-Drop Multiplexer (SONET ADM) Generic Criteria (a module of TSGR, FR-440)*, Issue 1 (Bellcore, December 1998).
- GR-499-CORE, *Transport System Generic Requirements (TSGR): Common Requirements (a module of TSGR, FR-440)*, Issue 2 (Bellcore, December 1998).
- GR-765-CORE, *Generic Requirements for Single Fiber Single-Mode Optical Splices and Splicing Systems*, Issue 1 (Bellcore, September 1995).
- GR-815-CORE, *Generic Requirements for Network Element/Network System (NE/NS) Security (a module of LSSGR, FR-64, and OTGR, FR-439)*, Issue 1 (Bellcore, November 1997).
- GR-820-CORE, *OTGR Section 5.1: Generic Digital Transmission Surveillance (a module of OTGR, FR-439)*, Issue 2 (Bellcore, December 1997).
- GR-826-CORE, *OTGR Section 10.2: User Interface Generic Requirements for Supporting Network Element Operations (a module of OTGR, FR-439)*, Issue 1 (Bellcore, June 1994).
- GR-828-CORE, *OTGR Section 11.2: Generic Operations Interface – OSI Communications Architecture (a module of OTGR, FR-439)*, Issue 1 (Bellcore, September 1994) plus Revision 2, October 1996.
- GR-831-CORE, *OTGR Section 12.1: Operations Application Messages - Language for Operations Application Messages (a module of OTGR, FR-439)*, Issue 1 (Bellcore, November 1996).

- GR-833-CORE, *OTGR Section 12.3: Network Maintenance: Network Element and Transport Surveillance Messages (a module of OTGR, FR-439)*, Issue 2 (Bellcore, November 1996).
- GR-834-CORE, *OTGR Section 12.4: Network Maintenance: Access and Testing Messages (a module of OTGR, FR-439)*, Issue 2 (Bellcore, November 1996).
- GR-836-CORE, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Overview: Transport Configuration and Surveillance for Network Elements (a module of OTGR, FR-439)*, Issue 2 (Bellcore, September 1996) plus Revisions.
- GR-836-IMD, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Details: Transport Configuration and Surveillance for Network Elements (a module of OTGR, FR-439)*, Issue 2 (Bellcore, September 1996).
- GR-839-CORE, *Generic Requirements for Supplier-Provided Training (a module of LSSGR, FR-64, TSGR, FR-440, and OTGR, FR-439)*, Issue 1 (Bellcore, July 1996).
- GR-910-CORE, *Generic Requirements for Fiber Optic Attenuators*, Issue 2 (Bellcore, December 1998).
- GR-1031-CORE, *OTGR Section 15.6: Operations Interfaces Using OSI Tools: Test Access Management*, Issue 2 (Bellcore, October 1997) plus Revision 1, December 1998.
- GR-1042-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Overview: Synchronous Optical Network (SONET) Transport Information Model*, Issue 3 (Bellcore, December 1998).
- GR-1042-IMD, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Details: Synchronous Optical Network (SONET) Transport Information Model*, Issue 3 (Bellcore, December 1998).
- GR-1093-CORE, *Generic State Requirements for Network Elements*, Issue 1 (Bellcore, October 1994) plus Revision 1, December 1995.
- GR-1209-CORE, *Generic Requirements for Fiber Optic Branching Components*, Issue 2 (Bellcore, February 1998).
- GR-1230-CORE, *SONET Bidirectional Line Switched Ring Equipment Generic Criteria (a module of TSGR, FR-440)*, Issue 4 (Bellcore, December 1998).
- GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria*, Issue 1 (Bellcore, June 1995).
- GR-1250-CORE, *Generic Requirements for Synchronous Optical Network (SONET) File Transfer*, Issue 1 (Bellcore, June 1995).

- GR-1253-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools: Telecommunications Management Network Security Administration*, Issue 1 (Bellcore, June 1995).
- GR-1309-CORE, *TSC/RTU and OTAU Generic Requirements for Remote Optical Fiber Testing (a module of OTGR, FR-439)*, Issue 1 (Bellcore, June 1995).
- GR-1332-CORE, *Generic Requirements for Data Communications Network Security*, Issue 2 (Bellcore, April 1996).
- GR-1377-CORE, *SONET OC-192 Transport System Generic Criteria (a module of TSGR, FR-440)*, Issue 5 (Bellcore, December 1998).
- GR-1400-CORE, *SONET Dual-Fed Unidirectional Path Switched Ring (UPSR) Equipment Generic Criteria (a module of TSGR, FR-440)*, Issue 2 (Bellcore, January 1999).
- GR-1469-CORE, *Generic Requirements on Security for OSI-Based Telecommunications Management Network Interfaces*, Issue 1 (Bellcore, September 1994).

Technical References

- TR-NWT-000057, *Functional Criteria for Digital Loop Carrier Systems (a module of TSGR, FR-440)*, Issue 2 (Bellcore, January 1993).
- TR-NWT-000078 (see GR-78-CORE).
- TR-NWT-000170, *Digital Cross-Connect System (DSC 1/0) Generic Criteria*, Issue 2 (Bellcore, January 1993).
- TR-NWT-000357, *Generic Requirements for Assuring the Reliability of Components Used in Telecommunication Systems (a module of RQGR, FR-796)*, Issue 2 (Bellcore, October 1993).
- TR-NWT-000418 (see GR-418-CORE).
- TR-TSY-000454 (see GR-454-CORE).
- TR-TSY-000458, *Digital Signal Zero, "A" (DS-0A 64 kb/s) Systems Interconnection*, Issue 1 (Bellcore, December 1989).
- TR-NWT-000468 (see GR-468-CORE).
- TR-TSY-000782, *SONET Digital Switch Trunk Interface Criteria (a module of LSSGR, FR-64, and TSGR, FR-440)*, Issue 2 (Bellcore, September 1989).
- TR-TSY-000824, *OTGR Section 10.1: User System Interface - User System Access (a module of OTGR, FR-439)*, Issue 2 (Bellcore, February 1988).

- TR-TSY-000825, *OTGR Section 10.A: User System Interface - User System Language* (a module of OTGR, FR-439), Issue 2 (Bellcore, February 1988).
- TR-TSY-000827, *OTGR Section 11.1: Generic Operations Interfaces: Non-OSI Communications Architecture* (a module of OTGR, FR-439), Issue 1 (Bellcore, November 1988).
- TR-NWT-000835, *OTGR Section 12.5: Network Element and Network System Security Administration Messages* (a module of OTGR, FR-439), Issue 3 (Bellcore, January 1993).
- TR-OPT-000839 (see GR-839-CORE).
- TR-NWT-000917, *SONET Regenerator (SONET RGTR) Equipment Generic Criteria* (a module of TSGR, FR-440), Issue 1 (Bellcore, December 1990).
- TR-NWT-000930, *Generic Requirements for Hybrid Microcircuits Used in Telecommunications Equipment* (a module of RQGR, FR-796), Issue 2 (Bellcore, September 1993).
- TR-NWT-001112, *Broadband-ISDN User to Network Interface and Network Node Interface Physical Layer Generic Criteria*, Issue 1 (Bellcore, June 1993).

Technical Advisories / Framework Technical Advisories

- TA-NPL-000286 (not available).
- TA-NPL-000464, *Generic Requirements and Design Considerations for Optical Digital Signal Cross-Connect Systems*, Issue 1 (Bellcore, September 1987).
- TA-NWT-000487 (see GR-487-CORE).
- TA-NWT-000983, *Reliability Assurance Practices for Optoelectronic Devices in Loop Applications*, Issue 2 (Bellcore, December 1993). (replaced by GR-468-CORE)
- TA-TSV-001294, *Generic Requirements for Element Management Layer (EML) Functionality and Architecture*, Issue 1 (Bellcore, December 1992).
- FA-NWT-001345, *Framework Generic Requirements for Element Manager (EM) Applications for SONET Subnetworks*, Issue 1 (Bellcore, September 1992).
- TA-NWT-001385 (not available).

Special Reports

- SR-104, *Bellcore Digest of Technical Information*, Volume 14, Issue 12 (Bellcore, December 1997).

- SR-NWT-002224 (not available).
- SR-TSV-002671, *EML Applications for Fault Management: Subnetwork Root Cause Alarm Analysis*, Issue 1 (Bellcore, June 1993).
- SR-TSV-002672, *EML Applications for Fault Management: Intelligent Alarm Filtering for SONET*, Issue 1 (Bellcore, March 1994).
- SR-TSV-002675, *EML Applications for Configuration Management: Resource Provisioning Selection and Assignment – Functional Description*, Issue 1 (Bellcore, December 1993).
- SR-TSV-002678, *EML Applications for Configuration Management: Inventory Notification and Query – Functional Description*, Issue 1 (Bellcore, April 1994).
- SR-NWT-002723, *Applicable TLI Messages for SONET Network Elements*, Issue 1 (Bellcore, June 1993).

NOTE:

All Bellcore documents are subject to change, and their citation in this document reflects the most current information available at the time of this printing. Readers are advised to check current status and availability of all documents. Documents can be ordered from Bellcore's on-line catalog or by contacting Bellcore.

To Contact Bellcore

Bellcore Customer Service
8 Corporate Place, Room 3A-184
Piscataway, New Jersey 08854-4156
1-800-521-CORE (2673) (USA and Canada)
1 (732) 699-5800 (all others)
1 (732) 336-2559 (FAX)

To Order Documents On-Line

- Perform the following steps to order from the on-line catalog:
 1. Enter the URL line: ***telecom-info.bellcore.com***
 2. Click on the ***Search*** button located on top
 3. In the ***Keywords*** field, enter the document number (or keywords), then click on ***Submit Search***.
- or . . .
1. Enter the above URL line
 2. Click on the ***Browse*** button located on top, then click on the subject of interest.

- Bellcore employees should perform the following steps:
 1. Access the Bellcore Internal Home Page
 2. Click on the *Marketwise* button located on top
 3. Click on *DOCS - Bellcore Product Catalog*
 4. Enter data in one or more of the fields (e.g., enter a document number in the *Product Number* field), then follow the instructions to search the on-line catalog.

EIA/TIA Documents

- EIA/TIA-455-170, *Cutoff Wavelength of Single-Mode Fiber by Transmitted Power.*
- EIA/TIA-492, *Generic Specification for Optical Waveguide Fiber.*
- EIA/TIA-559, *Single-Mode Fiber Optic System Transmission Design.*
- FOTP-127, *Spectral Characterization of Multimode Laser Diodes.*
- OFSTP-2, *Effective Transmitter Output Power Coupled into Single-Mode Fiber Optic Cable.*
- OFSTP-3, *Fiber Optic Terminal Receiver Sensitivity and Maximum Receiver Input Power.*
- OFSTP-10, *Measurement of Dispersion Power Penalty in Single-Mode Systems.*
- OFSTP-11, *Measurement of Single Reflection Power Penalty for Fiber Optic Terminal Equipment.*

These publications are available from:

EIA/TIA Standards Sales Office
2001 Pennsylvania, NW
Washington, DC 20006
(202) 457-4963

American National Standards Institute (ANSI) Documents

- ANSI T1.101-1994, *Synchronization Interface Standards for Digital Networks*
- ANSI T1.102-1993, *Digital Hierarchy-Electrical Interfaces*
- ANSI T1.105-1995, *Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates and Formats*
- ANSI T1.105.01-1995, *Synchronous Optical Network (SONET) - Automatic Protection Switching*
- ANSI T1.105.03-1994, *Synchronous Optical Network (SONET) - Jitter at Network Interfaces*
- ANSI T1.105.03a-1995, *Synchronous Optical Network (SONET) - Jitter at Network Interfaces - DS1 Supplement*
- ANSI T1.105.03b-1997, *Synchronous Optical Network (SONET) - Jitter at Network Interfaces - DS3 Supplement*
- ANSI T1.105.04-1995, *Synchronous Optical Network (SONET): Data Communication Channel Protocols and Architectures*

- ANSI T1.105.05-1994, *Synchronous Optical Network (SONET) - Tandem Connection Maintenance*
- dpANS T1.105.06-199x, *SONET Physical Layer Specifications*
- ANSI T1.105.09-1996, *SONET: Network Element Timing and Synchronization*
- ANSI T1.107-1995, *Digital Hierarchy - Formats Specifications*
- ANSI T1.204-1994, *Operations, Administration, Maintenance, and Provisioning-Lower Layer Protocols for Telecommunication Management Network (TMN) Interfaces Between Operations Systems and Network Elements*
- ANSI T1.210-1993, *Operations, Administration, Maintenance and Provisioning (OAM&P) — Principles of Functions, Architectures and Protocols for Telecommunications Management Network (TMN) Interfaces*
- ANSI T1.231-1997, *Digital Hierarchy - Layer 1 In-Service Digital Transmission Performance Monitoring*
- ANSI T1.245-1997, *Directory Service for Telecommunications Management Network (TMN) and Synchronous Optical Network (SONET)*
- ANSI T1.506A-1992, *Telecommunications - Network Performance - Specifications for Switched Exchange Network (Absolute Round-Trip Delay)*
- ANSI T1.508-1992, *Telecommunications - Network Performance - Loss Plan for Evolving Digital Networks*
- ANSI T1.508A-1993, *Telecommunications - Network Performance - Loss Plan for Evolving Digital Networks*
- ANSI T1.646-1995, *Broadband ISDN - Physical Layer Specification for User-Network Interfaces Including DS1/ATM*
- ANSI Technical Report #6, *A Technical Report on Slave Stratum Clock Performance Measurement Guidelines*
- T1 Technical Report #33, *A Technical Report on Synchronization Network Management Using Synchronization Status Messages*
- ANSI X3.216-1992, *Data Communications -Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address*

These publications are available from:

American National Standards Institute, Inc.
11 West 42nd Street
New York, NY 10036

U.S. Government Publications

U.S. Department of Health, Education, and Welfare; Bureau of Radiological Health.

- 21 CFR 1040.10, *Performance Standard for Laser Products*.

This publication is available from:

Director, Division of Compliance
Bureau of Radiological Health
5600 Fishers Lane
Rockville, MD 20857

ITU-T and CCITT Recommendations

- G.703, *Physical/electrical characteristics of hierarchical digital interfaces*
- G.707, *Network node interfaces for the synchronous digital hierarchy (SDH)*, March 1996
- G.709, *Synchronous multiplexing structure*
- G.783 (draft), *Characteristics of Synchronous Digital Hierarchy (SDH) equipment functional blocks*
- G.831, *Management capabilities of transport networks based on the Synchronous Digital Hierarchy (SDH)*, August 1996
- G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*
- G.958, *Digital line systems based on the synchronous digital hierarchy for use on optical fiber cables*
- M.3010, *Principles for a telecommunications management network*
- M.3100, *Generic Network Information Model*
- Q.811, *Lower Layer protocol profiles for the Q3 and X interfaces*
- Q.921, *ISDN user-network interface - Data link layer specification*
- X.121, *International numbering plan for public data networks*
- X.226, *Presentation Protocol Specification for Open Systems Interconnection for CCITT applications*.

These publications are available from:

International Telecommunication Union
General Secretariat — Sales Section

Place des Nations, CH-1211 Geneva 20 (Switzerland)
Tf: +41 22 730 5285

ISO Documents

- ISO/IEC 7498-1:1994, *Information technology - Open System Interconnection Reference Model: The Basic Model*
- ISO/IEC 8073:1988: *Information processing systems - Open Systems Interconnection - Connection oriented transport protocol specifications - Addendum 2: Class four operation over connectionless network service*
- ISO/IEC 8073:1988/Addendum 2:1989: *Information processing systems - Open Systems Interconnection - Connection oriented transport protocol specifications*
- ISO/IEC 8073:1992: *Information technology - Telecommunications and information exchange between systems - Open Systems Interconnection - Protocol for providing the connection-mode transport service*
- ISO/IEC 8208/CCITT X.25, *Information processing systems - X.25 Packet Level Protocol for Data Terminal Equipment*
- ISO/IEC 8327-1:1988: *Information processing systems - Open Systems Interconnection - Connection oriented session protocol specification - Part 1: Protocol specification*
- ISO/IEC DIS 8327-2: *Information technology - Open Systems Interconnection - Basic connection oriented session protocol specification*
- ISO 8348:1993, *Information processing systems - Data Communications - Network service definition*
- ISO/IEC 8473:1988: *Information processing systems - Data communications - Protocol for providing the connectionless-mode network layer service*
- ISO/IEC DIS 8473-1:1993: *Information technology - Protocol for providing the connectionless-mode network service*
- ISO 8648, *Internal Organization of the Network Layer*
- ISO/IEC 8650:1988: *Information technology - Open Systems Interconnection - Protocol specification for the Association Control Service Element*
- ISO/IEC DIS 8650-2: *Information technology - Open Systems Interconnection - Protocol specification for the Association Control Service Element*
- ISO 8802-2/ANSI/IEEE Std. 802.2-1989, *Information processing systems — Open Systems Interconnection — Local area networks — Part 2: Logical link control*

- ISO/IEC 8802-2:1989/Amd 3 - *Information processing systems - Local Area Networks - Part 2: Logical Link Control - Amendment 3: Conformance Requirements*
- ISO 8802-3/ANSI/IEEE Std. 802.3-1990, *Information processing systems — Open Systems Interconnection — Local area networks — Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*
- ISO/IEC 8823-1:1988: *Information processing systems - Open Systems Interconnection - Connection oriented presentation protocol specification - Part 1: Protocol specification*
- ISO/IEC DIS 8823-2: *Information technology - Open Systems Interconnection - Connection oriented presentation protocol specification*
- ISO/IEC 9542:1988: *Information processing systems - Telecommunications and information exchange between systems - End system to Intermediate system routing exchange protocol for use in conjunction with protocol for providing the connectionless-mode Network Service (ISO 8473)*
- ISO/IEC TR 9577, *Information processing systems - Data Communications - Protocol identification in the network layer*
- ISO/IEC 10040:1992: *Information technology - Open Systems Interconnection - Systems Management Overview*
- ISO TR 10172:1991: *Information technology - Telecommunications and information exchange between systems - Network/transport protocol interworking specification*
- ISO/IEC 10589:1992: *Information technology - Telecommunications and information exchange between systems - Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*
- ISO/IEC ISP 10607-1: *Information technology - International Standardized Profiles AFTnn - File Transfer, Access and Management*
- ISO/IEC ISP 10608-1:1992: *Information technology - International Standardized Profile TAnnnn - Connection-mode Transport Service over Connectionless-mode Network Service*
- ISO/IEC ISP 10608-2:1992: *Information technology - International Standardized Profile TAnnnn - Connection-mode Transport Service over Connectionless-mode Network Service*
- ISO DIS 10747:1992, *Information Technology — Telecommunications and Information Exchange between Systems-Protocol for Exchange of Inter-Domain Routing Information among Intermediate Systems to support forwarding of ISO/IEC 8473 PDUs*

- ISO/IEC ISP 11183-1: *Information technology - International Standardized Profiles AOM1n OSI Management- Management Communications*
- ISO/IEC PDISP 11188-1: *Information technology - International Standardized Profile - Common Upper Layer Requirements - Part 1: Basic connection oriented requirements.*

These publications are available from:

American National Standards Institute, Inc.
1430 Broadway
New York, NY 10018

IEEE Documents

- IEEE 802.2, *Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method*
- IEEE 802.3, *Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method*
- IEEE P802.6/D14, *Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN), September 1990.*

These publications can be obtained by calling:

IEEE Standards Publications
(800) 678-IEEE or (908) 981-1393

SONET Interoperability Forum (SIF) Documents

- SIF-002-1996, *Remote Login Implementation Requirements Specification, Issue 1*
- SIF-009-1997, *NE-NE Remote Login Implementation Requirements Specification*

The SIF is an Alliance for Telecommunications Industry Solutions (ATIS) sponsored committee. SIF-approved documents may be obtained directly from the ATIS WWW server at:

<http://www.atis.org/atis/sif/sifdoc.htm>

For additional information on SIF documents, contact:

Alliance for Telecommunications Industry Solutions
Attn.: Lisa Colaianne
1200 G Street, NW, Suite 500
Washington, DC 20005
Tel: 202-434-8823

Glossary

This section includes definitions and a list of acronyms.

Definitions

Add-Drop Multiplex (ADM) – A network element that provides access to some or all of the STS and/or VT paths contained within the OC-N optical signals at one or two OC-N interfaces. If two OC-N interfaces are provided, the path signals are added (inserted) to and/or dropped (extracted) from OC-N signals as they pass through the ADM.

Alarm Indication Signal (AIS) – A code sent downstream in a digital network to indicate that a traffic-related defect has been detected.

Asynchronous Transfer Mode (ATM) – A multiplexing/switching technique in which information is organized into fixed-length cells with each cell consisting of an identification header field and an information field. The transfer mode is asynchronous in the sense that the use of the cells depends on the required or instantaneous bit rate.

Bit Interleaved Parity N (BIP-N) – A method of error monitoring. If “even parity” is used, the transmitting equipment generates an N-bit code over a specified portion of the signal in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the covered portion of the signal, the second bit provides even parity over the second bits of all N-bit sequences within the specified portion, etc. Even parity is generated by setting the BIP-N bits so that there are an even number of ones in each of all N-bit sequences, including the BIP-N.

Concatenated Synchronous Transport Signal level N (STS-Nc) – A signal in which the STS Envelope Capacities from the N STS-1s have been combined to carry an STS-Nc Synchronous Payload Envelope (SPE). An STS-Nc may be transported as an OC-N or STS-N electrical signal, or it may be a module that is multiplexed into a higher rate signal (in which case it is referred to as an STS-Mc). In either case, it must be transported as a single entity, not as N (or M) separate signals.

Defect – A limited interruption in the ability of an item to perform a required function.

Digital Cross-Connect System (DCS) – A network element that terminates standard digital signals and facilities operating at a standard digital signal rate, and automatically cross-connects constituent (tributary) signals according to an electronically alterable memory map.

Distributed Queue Dual Bus (DQDB) – A multiplexing/switching technique similar to ATM that uses IEEE 802.6 PLCP.

Drop-side signal – A signal with reduced overhead functionality suitable for intra-office interconnection.

DS0 Path Terminating Equipment (DS0 PTE) – Network elements that multiplex/demultiplex the DS0 channels. DS0 PTEs interpret and either modify or create the DS0 signaling information necessary to transport the DS0 channels.

Electrical Carrier level 1 (EC-1) – One designation for the electrical interface signal that is the counterpart to the basic module in SONET, the STS-1. In this document, the term “STS-1 electrical” is used instead of “EC-1”.

Electrical Carrier level N (EC-N) – One designation for the electrical interface signal that is the counterpart to an STS-N. In this document, the term “STS-N electrical” is used instead of “EC-N”.

Eye Diagram – A graphic presentation formed by the superimposition of the waveforms of all possible pulse sequences.

Failure – A termination of the ability of an item to perform a required function. A failure is caused by the persistence of a defect.

Far End Block Error (FEBE) – See Remote Error Indication (REI).

Fixed Stuff (R-Bits/Bytes) – Fixed stuff (R) bits and bytes are used to compensate for the differences between the bandwidth available in the STS-1 and VT Synchronous Payload Envelopes and the bandwidth required for the actual payload mappings (e.g., DS1, DS1C, DS2, and DS3). R-bits and bytes have no defined value. The receiver is required to ignore the value of these bits/bytes (except for BIP-8 calculation/verification).

Gateway Communications Functions – Functions to facilitate operations communications between two communicating entities across dissimilar subnetworks. Examples include concentration, message routing and relaying (Network Layer), and application layer protocol conversion and/or message translation. See Mediation Functions.

Information Processing Functions – Functions centralized in a subnetwork to perform common application processing and management for NEs within the subnetwork. Example NE functions include performance data storage and failure condition threshold crossing detection for alarm reporting. Examples of management functions include the migration of “OS-like” functions out to the network for subnetwork trouble sectionalization, configuration (e.g., cross-connection) management, and Customer Network Management (CNM). Also see Mediation Functions.

Inter-Carrier Interface (ICI) – The interface between two networks that belong to different network providers or carriers.

Line – A transmission medium, together with the associated Line Terminating Equipment (LTE), required to provide the means of transporting information between two consecutive line terminating network elements, one of which originates the line signal and the other terminates the line signal. See Figures 2-1 and 2-2.

Line Alarm Indication Signal (AIS-L) – AIS-L is generated by Section Terminating Equipment (STE) upon the detection of an Loss of Signal or Loss of Frame defect, or an equipment failure. AIS-L maintains operation of the downstream regenerators, and therefore prevents generation of unnecessary alarms. At the same time, data and orderwire communication is retained between the regenerators and the downstream Line Terminating Equipment (LTE).

Line Remote Defect Indication (RDI-L) – A signal returned to the transmitting Line Terminating Equipment (LTE) upon detecting a Loss of Signal, Loss of Frame, or AIS-L defect. RDI-L was previously known as Line FERF.

Line-Side Signal – A signal with full overhead functionality suitable for inter-office interconnection.

Line Terminating Equipment (LTE) – Equipment that terminates the SONET Line layer. LTE interprets and modifies or creates the Line Overhead. An NE that contains LTE will also contain STE.

Mediation Functions – Usually consist of Gateway Communications Functions, but can also include Information Processing Functions. When the two functions are combined, they are often contained in a stand-alone Mediation Device (MD), or packaged as an added module to an NE or equipment frame. Gateway Communications Functions often exist alone in a Gateway NE or Intermediate NE.

Most-Significant Bit – The left-most bit position, Bit 1, as illustrated in Figure 3-2. In SONET, the most-significant bit is transmitted first.

Optical Carrier level 1 (OC-1) – The optical interface signal that is the counterpart to the basic module in SONET, the STS-1.

Optical Carrier level N (OC-N) – The optical interface signal that is the counterpart to an STS-N.

Path – A path at a given bit rate is a logical connection between the point at which a standard frame format for the signal at the given bit rate is assembled, and the point at which the standard frame format for the signal is disassembled. See 2-1 and 2-2.

Path Overhead (POH) – Overhead assigned to and transported with the payload until the payload is demultiplexed. It is used for functions that are necessary to transport the payload.

Payload Pointer – The pointer that indicates the location of the beginning of the Synchronous Payload Envelope.

Pigtail – A length of optical fiber with one end terminated at a connector and the other end attached to a light source or detector. It is used to couple light from a source to a connectorized fiber cable or from a fiber cable to a detector.

Remote Alarm Indication (RAI) – A code sent upstream in a DS_n network as a notification that a failure condition has been declared downstream. (RAI signals were previously referred to as Yellow signals.)

Remote Error Indication (REI) – An indication returned to a transmitting node (source) that an errored block has been detected at the receiving node (sink). This indication was formerly known as Far End Block Error (FEBE).

RGTR (SONET Regenerator) – A unidirectional device that can receive a digital signal and retransmit it in a form in which the amplitude, waveforms, and timing characteristics of the signal are constrained within specified limits. In general, the SONET RGTR is defined as Section Terminating Equipment.

Section – The portion of a transmission facility, including terminating points, between (i) a terminal network element and a regenerator or (ii) two regenerators. A terminating point is the point after signal regeneration at which performance monitoring is (or may be) done. See Figures 2-1 and 2-2.

Section Terminating Equipment (STE) – Equipment that terminates the SONET Section layer. STE interprets and modifies or creates the Section Overhead.

SONET – An acronym for Synchronous Optical NETwork. SONET is a term in general usage that refers to the rates and formats that this standard specifies.

SONET Transport Overhead – The overhead added to the STS SPE for transport purposes. Transport Overhead consists of Line and Section Overhead. See Section 3.2.

STS Envelope Capacity – Bandwidth within, and aligned to, the STS Frame that carries the STS Synchronous Payload Envelope (SPE). The bandwidth from N STS-1s can be combined to carry an STS-Nc SPE.

STS Path Overhead (STS POH) – Nine evenly distributed Path Overhead bytes per 125 μ s starting at the first byte of the STS SPE. STS POH provides for communication between the point of creation of an STS SPE and its point of disassembly. STS POH is described in Section 3.3.2.3.

STS Path Remote Defect Indication (RDI-P) – A signal returned to the transmitting STS Path Terminating Equipment (PTE) upon detection of certain defects on the incoming path.

STS Path Terminating Equipment (STS PTE) – Equipment that terminates the SONET STS Path layer. STS PTE interprets and modifies or creates the STS Path Overhead. An NE that contains STS PTE will also contain LTE and STE.

STS-1 Electrical Signal – The electrical interface signal that is the counterpart to the basic module in SONET, the STS-1. An STS-1 electrical signal may also be referred to as an EC-1 signal.

STS-N Electrical Signal – The electrical interface signal that is the counterpart to an STS-N. An STS-N electrical signal may also be referred to as an EC-N signal.

STS-N Tandem Connection – A group of N STS-1s that are transported and maintained together through one or more tandem line systems, with the constituent SPE payload capacities unaltered. The tandem connection sub-layer falls between the SONET Line and STS Path layers.

Super Rate Payload – A signal that has to be carried by a Concatenated Synchronous Transport Signal level N (STS-Nc).

Synchronous – The essential characteristic of time scales or signals such that their corresponding significant instants occur at precisely the same average rate.

Synchronous Network – The synchronization of synchronous transmission systems with synchronous payloads to a master (network) clock that can be traced to a reference clock.

Synchronous Payloads – Payloads derivable from a network transmission signal by removing integral numbers of bits in every frame (i.e., there are no variable bit stuffing rate adjustments required to fit the payload in the transmission signal).

Synchronous Transport Signal level 1 (STS-1) – The basic (functional) module used to build SONET signals. An STS-1 has a bit rate of 51.84 Mb/s, and may be converted to an OC-1 or STS-1 electrical interface signal, multiplexed with other modules to form a higher rate (STS-N) signal, or combined with other STS-1s to form an STS-Nc.

Synchronous Transport Signal level N (STS-N) – A (functional) module used to build SONET signals. An STS-N has a bit rate of $N \times 51.84$ Mb/s, and may be converted to an OC-N or STS-N electrical interface signal, or multiplexed with other modules to form a higher rate signal (in which case it is referred to as an STS-M).

Terminal Multiplex (TM) – A network element that provides access to all of the STS and/or VT paths contained within the OC-N optical signals at one OC-N interface. An ADM with only one OC-N interface may be referred to as a TM.

Undefined bits/bytes – Those locations within the signal that do not have a function or value assigned to them. The receiver is required to ignore the value of these bits/bytes (except for BIP-8 calculation/verification).

Unequipped Channel – A portion of an STS-N such as an STS-1 SPE or VT SPE that is intentionally unoccupied.

Unequipped Indication – A code that originating equipment places in unequipped channels to indicate to Path Terminating Equipment that the channel is intentionally unoccupied.

User Channel – A channel that is allocated to the user (unless stated otherwise, “user” in this document refers to the network provider) for input of information such as data communication for use in maintenance activities and remoting of alarms external to the span equipment in a proprietary fashion.

Virtual Tributary (VT) – A structure designed for transport and switching of sub-STS-1 payloads. There are currently four sizes of VT as defined in Section 3.2.4.

VT Envelope Capacity – Bandwidth within, and aligned to, the VT Superframe that is available for the VT Synchronous Payload Envelope.

VT Group – A 9-row by 12-column structure (108 bytes) that carries one or more VTs of the same size. Seven VT groups (756 bytes) are byte interleaved within the VT-structured STS-1 SPE.

VT Path Overhead (V5, J2, Z6, and Z7) – Four evenly distributed path overhead bytes per VT SPE, starting with the first byte of the VT SPE. VT Path Overhead provides for communication between the point of creation of a VT SPE and its point of disassembly. VT Path Overhead is described in Section 3.3.3.

VT Path Remote Defect Indication (RDI-V) – A signal returned to the transmitting VT PTE upon detection of certain defects on the incoming path.

VT Path Remote Failure Indication (RFI-V) – A signal, applicable only to a VT1.5 with the byte-synchronous DS1 mapping, that is returned to the transmitting VT PTE upon declaring certain failures. The RFI-V signal was previously known as the VT Path Yellow signal.

VT Path Terminating Equipment (VT PTE) – Equipment that terminates the SONET VT Path layer. VT PTE interprets and modifies or creates the VT Path Overhead. An NE that contains VT PTE will also contain STS PTE, LTE, and STE.

VT Payload Capacity – The maximum bandwidth within the VT Synchronous Payload Envelope that is available for payload.

VT Superframe – The VT is organized into a 500- μ s superframe structure overlaid on and aligned to the 125- μ s STS-1 Synchronous Payload Envelope (SPE). The VT Payload Pointer and the VT SPE are contained within this structure.

VT Synchronous Payload Envelope (VT SPE) – A 500- μ s frame structure carried by the VT and composed of VT Path Overhead and bandwidth for payload. The envelope is contained within, and can have any alignment with respect to, the VT Envelope Capacity. The term generically refers to VT1.5, VT2, VT3, and VT6 SPEs, which are described in Section 3.2.4.

VT_x – A VT of size “x” (x = 1.5, 2, 3, or 6).

Work Station (WS) – Any one of a variety of Visual Display Terminals (VDTs), ranging from a simple keyboard and display to an intelligent, processor-controlled VDT.

Yellow Signal – See Remote Alarm Indication (REI) and VT Path Remote Failure Indication (RFI-V).

Acronyms

ACI	Application Context Identifier
ACSE	Association Control Service Element
ADM	Add-Drop Multiplex
AFI	Authority and Format Identifier
AID	Access Identifier
AIS	Alarm Indication Signal
AITs	Acknowledged Information Transfer Service
ANSI	American National Standards Institute
APD	Avalanche Photodiode
APS	Automatic Protection Switching
ASE	Application Service Element
ASN.1	Abstract Syntax Notation 1
ATM	Asynchronous Transfer Mode
B3ZS	Bipolar with Three-Zero Substitution
B8ZS	Bipolar with Eight-Zero Substitution
BCC	Bellcore Client Company
BER	Bit Error Ratio
BIP	Bit Interleaved Parity
B-ISDN	Broadband Integrated Services Digital Network
BITS	Building Integrated Timing Supply
CC	Composite Clock
CCITT	International Telegraph & Telephone Consultative Committee (replaced by ITU-T)
CEV	Controlled Environmental Vault
CID	Calling Address Identification
CLNP	Connectionless-mode Network Layer Protocol
CLNS	Connectionless-mode Network Service
CMI	Coded Mark Inversion
CMISE	Common Management Information Service Element
CNM	Customer Network Management
CO	Central Office
CONP	Connection-mode Network Layer Protocol
CR	Critical alarm
CSMA/CD	Carrier Sense Multiple Access with Collision Detection
CV	Coding Violation
CV-L	Line Coding Violation

CV-LFE	Far End Line Coding Violation
CV-P	STS Path Coding Violation
CV-PFE	Far End STS Path Coding Violation
CV-S	Section Coding Violation
CV-V	VT Path Coding Violation
CV-VFE	Far End VT Path Coding Violation
DCC	Data Communications Channel
DCN	Data Communications Network
DCS	Digital Cross-Connect System
DFB	Distributed Feedback
DM	Degraded Minute
dpANS	Draft Proposed American National Standard
DQDB	Distributed Queue Dual Bus
DSP	Domain Specific Part
DS	Digital Signal
DSn	Digital Signal at level n
DSNE	Directory Server NE
EIA	Electronic Industries Association
ENE	End NE
EOC	Embedded Operations Channel
EOW	Express Orderwire
ES	End System
ES	Errored Second
ES-L	Line Errored Second
ES-LFE	Far End Line Errored Second
ES-P	STS Path Errored Second
ES-PFE	Far End STS Path Errored Second
ES-S	Section Errored Second
ES-V	VT Path Errored Second
ES-VFE	Far End VT Path Errored Second
ESF	Extended Superframe Format
FA	Framework Technical Advisory
FC	Failure Count
FC-L	Line Failure Count
FC-LFE	Far End Line Failure Count
FC-P	STS Path Failure Count
FC-PFE	Far End STS Path Failure Count
FC-V	VT Path Failure Count

FC-VFE	Far End VT Path Failure Count
FDDI	Fiber Distributed Data Interface
FEBE	Far End Block Error (replaced with REI)
FERF	Far End Receive Failure (replaced with RDI)
FR	Family of Requirements
FT	File Transfer
FTAM	File Transfer Access and Management
GNE	Gateway NE
GOSIP	U.S. Government OSI Profile
GR	Generic Requirement
IAO	Intraoffice
ICI	Inter-Carrier Interface
ID	Identifier
IDI	Initial Domain Identifier
IDLC	Integrated Digital Loop Carrier
IDP	Initial Domain Part
IDRP	Inter Domain Routing Protocol
IEEE	Institute of Electrical and Electronics Engineers
INE	Intermediate NE
IR	Intermediate Reach
IS	Intermediate System
ISI	Intersymbol Interference
ISO	International Organization for Standardization
ISP	International Standardized Profile
ITU-T	International Telecommunication Union – Telecommunication Standardization Sector (formerly CCITT)
LAN	Local Area Network
LAPD	Link Access Protocol on the D-Channel
LAPD+	Link Access Protocol for non-D-Channel (applications)
LBC	Laser Bias Current
LBO	Line Build Out
LCN	Local Communications Network
LDB	Loop Detection Buffer
LEC	Local Exchange Carrier
LED	Light-Emitting Diode
LOF	Loss Of Frame
LOP	Loss Of Pointer
LOS	Loss Of Signal

LOW	Local Orderwire
LR	Long Reach
LSS	Link Status Signal
LTE	Line Terminating Equipment
MAC	Media Access Control
MAN	Metropolitan Area Network
MD	Mediation Device
MF	Mediation Function
MJ	Major alarm
MLM	Multi-Longitudinal Mode
MN	Minor alarm
MPN	Mode Partition Noise
MTIE	Maximum Time Interval Error
MTTR	Mean Time To Repair
NA	Not Alarmed
NDF	New Data Flag
NE	Network Element
NET	Network Entity Title
NPDU	Network Protocol Data Unit
NRZ	Non-Return to Zero
NSA	Non-Service Affecting
NSAP	Network Service Access Point
NSPDU	Network Service Protocol Data Unit
OAM&P	Operations, Administration, Maintenance, & Provisioning
OC-N	Optical Carrier at level N
OIW	OSI Implementors Workshop
OOF	Out Of Frame
OPR	Optical Power Received
OPT	Optical Power Transmitted
ORL	Optical Return Loss
OS	Operations System
OSE	Open Systems Environment
OSI	Open Systems Interconnection
OTGR	Operations Technology Generic Requirements
OW	Orderwire
PBX	Private Branch Exchange
PDU	Protocol Data Unit
PID	Protocol Identification

PIN	Positive-Intrinsic-Negative (photodiode)
PJ	Pointer Justification
PLCP	Physical Layer Convergence Procedure
PLM-P	STS Path Payload Label Mismatch
PLM-V	VT Path Payload Label Mismatch
PM	Performance Monitoring
PPDU	Presentation Protocol Data Unit
PSAP	Presentation Service Access Point
PSC	Protection Switching Count
PSD	Protection Switching Duration
PSN	Packet Switched Network
PTE	Path Terminating Equipment
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RAI	Remote Alarm Indication
RDI	Remote Defect Indication
REI	Remote Error Indication
RFI	Remote Failure Indication
RGTR	Regenerator
RPDU	Route Protocol Data Unit
RPP	Reliability Prediction Procedure
RZ	Return to Zero
SA	Service Affecting
SAPI	Service Access Point Identifier
SD	Signal Degrade
SDH	Synchronous Digital Hierarchy
SEF	Severely Errored Framing
SEFS	Severely Errored Framing Second
SEFS-S	Section Severely Errored Framing Second
SES	Severely Errored Second
SES-L	Line Severely Errored Second
SES-LFE	Far End Line Severely Errored Second
SES-P	STS Path Severely Errored Second
SES-PFE	Far End STS Path Severely Errored Second
SES-S	Section Severely Errored Second
SES-V	VT Path Severely Errored Second
SES-VFE	Far End VT Path Severely Errored Second
SF	Signal Fail

SIA	Stable Implementation Agreements
SLM	Single Longitudinal Mode
SLM-P	STS Signal Label Mismatch (replaced by PLM-P and UNEQ-P)
SLM-V	VT Signal Label Mismatch (replaced by PLM-V and UNEQ-V)
SMC	SONET Minimum Clock
SMF	Single Mode Fiber
SNDCF	Subnetwork Dependent Convergence Function
SNR	Signal to Noise Ratio
SONET	Synchronous Optical Network
SPE	Synchronous Payload Envelope
SR	Short Reach
SR	Special Report
SSR	Side Mode Suppression Ratio
STE	Section Terminating Equipment
STS	Synchronous Transport Signal
STS-N	Synchronous Transport Signal level N
SYNTRAN	Synchronous (DS3) Transmission
TA	Technical Advisory
TARP	TID Address Resolution Protocol
TBD	To Be Determined
TCA	Threshold Crossing Alert
TDEV	Time Deviation
TDC	TARP Data Cache
TEF	TARP Echo Function
TEI	Terminal Endpoint Identifier
TIA	Telecommunications Industries Association
TID	Target Identification
TIE	Time Interval Error
TL1	Transaction Language 1
TM	Terminal Multiplex
TMN	Telecommunications Management Network
TP	Transport Protocol
TP4	Transport Protocol Class 4
TPF2	Transaction Processing Full conformance stack #2
TR	Technical Reference
TSAP	Transport Service Access Point
UAS	Unavailable Second
UAS-L	Line Unavailable Second

UAS-LFE	Far End Line Unavailable Second
UAS-P	STS Path Unavailable Second
UAS-PFE	Far End STS Path Unavailable Second
UAS-V	VT Path Unavailable Second
UAS-VFE	Far End VT Path Unavailable Second
UI	Unit Interval
UID	User Identification
UITS	Unacknowledged Information Transfer Service
UNEQ-P	STS Path Unequipped
UNEQ-V	VT Path Unequipped
UNI	User Network Interface
URC	Update Remote Cache
USL	User System Language
VC	Virtual Circuit
VT	Virtual Terminal
VT	Virtual Tributary
VTE	Virtual Terminal Environment
WS	Workstation
WTR	Wait To Restore
ZBTSI	Zero Byte Time Slot Interchange
ZCS	Zero Code Suppression

Requirement-Object Index

[1]	CR2-1	Page 2-3	[39]	R3-36	Page 3-46
[2]	CR2-2	Page 2-3	[40]	CR3-37	Page 3-47
[3]	R2-3	Page 2-4	[41]	CR3-38	Page 3-47
[4]	R3-1	Page 3-2	[42]	R3-39	Page 3-47
[5]	O3-2	Page 3-2	[43]	R3-40	Page 3-47
[6]	R3-3	Page 3-3	[44]	CR3-41	Page 3-47
[7]	R3-4	Page 3-3	[45]	CR3-42	Page 3-48
[8]	R3-5	Page 3-4	[46]	R3-43	Page 3-48
[9]	R3-6	Page 3-4	[47]	R3-44	Page 3-48
[10]	R3-7	Page 3-5	[48]	CR3-45	Page 3-49
[11]	R3-8	Page 3-7	[49]	CR3-46	Page 3-49
[12]	R3-9	Page 3-8	[50]	CR3-47	Page 3-49
[13]	R3-10	Page 3-10	[51]	R3-48	Page 3-49
[14]	R3-11	Page 3-10	[52]	R3-49	Page 3-49
[15]	R3-12	Page 3-10	[53]	R3-50	Page 3-49
[16]	R3-13	Page 3-29	[54]	R3-51	Page 3-49
[17]	O3-14	Page 3-29	[55]	R3-52	Page 3-50
[18]	R3-15	Page 3-29	[56]	R3-53	Page 3-50
[19]	R3-16	Page 3-30	[57]	R3-55	Page 3-50
[20]	R3-17	Page 3-31	[58]	R3-56	Page 3-50
[21]	R3-18	Page 3-31	[59]	R3-57	Page 3-51
[22]	R3-19	Page 3-32	[60]	R3-58	Page 3-51
[23]	R3-20	Page 3-33	[61]	R3-59	Page 3-51
[24]	R3-21	Page 3-33	[62]	R3-60	Page 3-51
[25]	R3-22	Page 3-34	[63]	R3-62	Page 3-51
[26]	R3-23	Page 3-34	[64]	R3-63	Page 3-52
[27]	R3-24	Page 3-36	[65]	R3-64	Page 3-52
[28]	R3-25	Page 3-39	[66]	R3-65	Page 3-53
[29]	R3-26	Page 3-39	[67]	R3-66	Page 3-53
[30v2]	R3-27	Page 3-39	[68]	R3-67	Page 3-53
[31]	R3-28	Page 3-39	[69]	R3-69	Page 3-55
[32]	R3-29	Page 3-41	[70]	R3-70	Page 3-55
[33]	R3-30	Page 3-42	[71]	R3-71	Page 3-55
[34]	R3-31	Page 3-43	[72]	R3-74	Page 3-56
[35]	R3-32	Page 3-45	[73]	R3-77	Page 3-58
[36]	R3-33	Page 3-45	[74]	R3-78	Page 3-59
[37]	O3-34	Page 3-45	[75]	R3-79	Page 3-59
[38]	CR3-35	Page 3-46	[76]	R3-80	Page 3-61
			[77]	R3-81	Page 3-61
			[78]	R3-82	Page 3-61
			[79]	R3-83	Page 3-63

[80]	R3-84	Page 3-63	[121]	R4-12	Page 4-12
[81]	R3-85	Page 3-63	[122]	O4-13	Page 4-12
[82]	R3-86	Page 3-64	[123]	R4-14	Page 4-23
[83]	R3-87	Page 3-64	[124]	CR4-15	Page 4-23
[84]	R3-88	Page 3-65	[125]	R4-16	Page 4-23
[85]	R3-89	Page 3-67	[126]	R4-17	Page 4-23
[86]	R3-90	Page 3-69	[127]	R4-18	Page 4-23
[87]	R3-91	Page 3-69	[128v2]	CR4-19	Page 4-43
[88]	R3-92	Page 3-69	[129]	R4-21	Page 4-43
[89]	R3-93	Page 3-69	[130]	R4-22	Page 4-47
[90]	O3-94	Page 3-69	[131]	R5-1	Page 5-1
[91]	R3-95	Page 3-72	[132]	R5-2	Page 5-2
[92]	R3-96	Page 3-72	[133]	R5-3	Page 5-2
[93]	R3-97	Page 3-72	[134]	R5-6	Page 5-6
[94]	R3-98	Page 3-72	[135]	R5-7	Page 5-10
[95]	R3-99	Page 3-72	[136]	CR5-9	Page 5-14
[96v2]	R3-100	Page 3-73	[137]	O5-10	Page 5-14
[97]	R3-102	Page 3-74	[138]	O5-11	Page 5-14
[98]	R3-103	Page 3-76	[139]	R5-12	Page 5-14
[99]	R3-104	Page 3-77	[140]	R5-13	Page 5-15
[100]	R3-105	Page 3-77	[141]	R5-14	Page 5-15
[101]	R3-106	Page 3-77	[142]	R5-15	Page 5-15
[102]	R3-107	Page 3-77	[143]	CR5-16	Page 5-15
[103]	O3-108	Page 3-77	[144]	CR5-17	Page 5-15
[104]	R3-109	Page 3-77	[145]	CR5-18	Page 5-16
[105]	R3-110	Page 3-78	[146]	R5-19	Page 5-16
[106]	R3-111	Page 3-78	[147]	CR5-20	Page 5-16
[107]	R3-112	Page 3-78	[148]	CR5-21	Page 5-16
[108v2]	R3-113	Page 3-79	[149]	CR5-22	Page 5-16
[109]	R3-115	Page 3-80	[150]	CR5-23	Page 5-17
[110]	R4-1	Page 4-4	[151]	R5-24	Page 5-17
[111]	R4-2	Page 4-4	[152]	R5-25	Page 5-17
[112]	R4-3	Page 4-5	[153]	CR5-26	Page 5-18
[113]	R4-4	Page 4-5	[154]	R5-27	Page 5-18
[114]	O4-5	Page 4-6	[155]	CR5-28	Page 5-19
[115]	R4-6	Page 4-6	[156]	CR5-29	Page 5-19
[116]	R4-7	Page 4-6	[157]	CR5-30	Page 5-19
[117]	R4-8	Page 4-6	[158]	CR5-31	Page 5-19
[118]	O4-9	Page 4-6	[159]	CR5-32	Page 5-19
[119]	R4-10	Page 4-8	[160]	R5-33	Page 5-19
[120]	R4-11	Page 4-12	[161]	O5-34	Page 5-20

[162]	R5-35	Page 5-20	[203]	R5-82	Page 5-39
[163]	CR5-36	Page 5-20	[204]	R5-83	Page 5-39
[164]	R5-37	Page 5-20	[205]	R5-84	Page 5-39
[165]	R5-38	Page 5-21	[206]	R5-85	Page 5-39
[166]	CR5-39	Page 5-21	[207]	R5-87	Page 5-40
[167]	R5-40	Page 5-21	[208]	O5-88	Page 5-40
[168]	R5-41	Page 5-21	[209]	R5-89	Page 5-41
[169]	R5-42	Page 5-21	[210]	R5-90	Page 5-41
[170]	O5-43	Page 5-21	[211]	R5-91	Page 5-42
[171]	R5-44	Page 5-22	[212]	R5-92	Page 5-42
[172]	O5-45	Page 5-22	[213]	CR5-93	Page 5-43
[173v2]	R5-46	Page 5-22	[214]	R5-94	Page 5-43
[174]	R5-47	Page 5-23	[215]	R5-95	Page 5-43
[175]	R5-48	Page 5-23	[216]	R5-96	Page 5-43
[176]	R5-50	Page 5-25	[217v2]	R5-97	Page 5-43
[177v3]	O5-53	Page 5-25	[222]	R5-104	Page 5-56
[178v2]	R5-57	Page 5-28	[223]	R5-105	Page 5-56
[179]	R5-58	Page 5-29	[224v2]	R5-99	Page 5-54
[180]	R5-59	Page 5-29	[225v3]	R5-100	Page 5-54
[181]	R5-60	Page 5-31	[226]	R5-106	Page 5-56
[182]	R5-61	Page 5-31	[227]	R5-107	Page 5-57
[183]	R5-62	Page 5-32	[239]	CR5-111	Page 5-58
[184]	R5-63	Page 5-33	[240v2]	R5-112	Page 5-60
[185]	R5-64	Page 5-33	[241]	R5-113	Page 5-60
[186]	R5-65	Page 5-33	[242]	R5-114	Page 5-60
[187]	R5-66	Page 5-33	[243v2]	CR5-115	Page 5-61
[188]	R5-67	Page 5-33	[244v2]	CR5-116	Page 5-61
[189]	R5-68	Page 5-34	[245v2]	CR5-117	Page 5-61
[190v2]	R5-69	Page 5-34	[246]	R5-118	Page 5-61
[191]	R5-70	Page 5-34	[247]	R5-119	Page 5-61
[192]	R5-71	Page 5-35	[248]	R5-121	Page 5-62
[193]	R5-72	Page 5-35	[249v3]	R5-122	Page 5-62
[194]	R5-73	Page 5-36	[250]	R5-127	Page 5-63
[195]	R5-74	Page 5-36	[253]	R5-128	Page 5-64
[196]	R5-75	Page 5-36	[254]	R5-129	Page 5-64
[197v2]	R5-76	Page 5-36	[255v2]	R5-131	Page 5-64
[198]	R5-77	Page 5-36	[256v2]	R5-132	Page 5-65
[199]	R5-78	Page 5-36	[257]	R5-134	Page 5-67
[200]	R5-79	Page 5-38	[258]	R5-135	Page 5-67
[201]	R5-80	Page 5-38	[259v2]	R5-136	Page 5-70
[202]	R5-81	Page 5-38	[260]	O5-138	Page 5-70

[261v2]	R5-140	Page 5-72	[314]	R5-199	Page 5-91
[262v2]	O5-141	Page 5-72	[315v2]	R5-200	Page 5-92
[264v2]	R5-133	Page 5-67	[316]	R5-201	Page 5-92
[265]	R5-142	Page 5-72	[317v2]	R5-203	Page 5-92
[266]	CR5-143	Page 5-72	[318]	R5-202	Page 5-92
[267v2]	R5-144	Page 5-72	[319]	R5-204	Page 5-92
[268]	R5-145	Page 5-72	[320v3]	O5-205	Page 5-93
[271v2]	R5-147	Page 5-73	[321]	R5-206	Page 5-93
[272]	R5-148	Page 5-73	[322]	R5-207	Page 5-93
[273v3]	R5-152	Page 5-74	[323]	R5-208	Page 5-93
[274v2]	CR5-153	Page 5-76	[324]	R5-209	Page 5-93
[276]	O5-160	Page 5-78	[325]	R5-210	Page 5-94
[277]	CR5-161	Page 5-78	[326v2]	R5-211	Page 5-94
[278]	R5-162	Page 5-78	[327v3]	R5-212	Page 5-94
[279v2]	R5-163	Page 5-78	[328]	R5-217	Page 5-96
[280v2]	R5-164	Page 5-78	[329]	R5-218	Page 5-96
[281]	R5-165	Page 5-78	[330]	R5-219	Page 5-96
[282]	R5-166	Page 5-78	[331]	R5-220	Page 5-97
[283]	R5-167	Page 5-79	[332]	R5-221	Page 5-97
[284v2]	R5-168	Page 5-79	[333]	R5-222	Page 5-98
[285v3]	R5-169	Page 5-81	[334]	R5-223	Page 5-98
[286]	R5-170	Page 5-82	[335]	R5-224	Page 5-98
[287v2]	CR5-171	Page 5-82	[336]	R5-225	Page 5-98
[288]	R5-172	Page 5-82	[337]	R5-226	Page 5-99
[289]	R5-173	Page 5-83	[338]	R5-227	Page 5-103
[290v2]	R5-174	Page 5-83	[339]	R5-228	Page 5-103
[291]	R5-175	Page 5-83	[340]	R5-229	Page 5-103
[292]	R5-176	Page 5-83	[341]	R5-230	Page 5-104
[293v2]	R5-177	Page 5-83	[342]	R5-231	Page 5-105
[294]	R5-178	Page 5-83	[343v2]	R5-232	Page 5-105
[295]	CR5-179	Page 5-83	[345]	R5-234	Page 5-106
[296]	R5-180	Page 5-84	[346]	R5-235	Page 5-108
[297]	R5-181	Page 5-84	[347]	R5-236	Page 5-108
[298v3]	R5-182	Page 5-84	[348]	R5-237	Page 5-108
[300v4]	R5-183	Page 5-85	[349]	O5-238	Page 5-109
[302v2]	R5-192	Page 5-88	[350]	R5-241	Page 5-111
[305]	R5-193	Page 5-88	[351]	R5-242	Page 5-112
[310]	R5-194	Page 5-89	[352]	R5-243	Page 5-112
[311]	R5-195	Page 5-89	[353]	R5-244	Page 5-113
[312]	R5-197	Page 5-91	[354]	O5-245	Page 5-113
[313]	R5-198	Page 5-91	[355]	R5-246	Page 5-114

[356]	O5-247	Page 5-114	[403]	R6-39	Page 6-11
[357]	R5-248	Page 5-114	[404]	R6-40	Page 6-11
[358]	R5-254	Page 5-119	[405]	R6-41	Page 6-12
[359]	R5-256	Page 5-122	[406]	R6-42	Page 6-12
[360v3]	R5-260	Page 5-126	[407]	R6-43	Page 6-12
[361v2]	R6-2	Page 6-2	[408]	R6-44	Page 6-12
[366]	R6-1	Page 6-2	[409]	R6-45	Page 6-12
[367]	R6-5	Page 6-4	[410]	R6-46	Page 6-12
[368]	R6-6	Page 6-4	[411]	R6-47	Page 6-12
[369]	R6-7	Page 6-5	[412]	O6-48	Page 6-12
[370]	R6-8	Page 6-5	[413]	R6-49	Page 6-13
[371]	R6-9	Page 6-5	[414]	O6-50	Page 6-13
[372]	O6-10	Page 6-5	[415]	O6-51	Page 6-13
[373]	R6-11	Page 6-5	[416v2]	R6-52	Page 6-16
[374]	R6-12	Page 6-5	[417]	R6-54	Page 6-16
[375]	R6-13	Page 6-5	[418v2]	R6-55	Page 6-16
[376]	R6-14	Page 6-5	[419]	R6-56	Page 6-16
[377]	CR6-15	Page 6-6	[420]	R6-57	Page 6-17
[378]	R6-16	Page 6-6	[421v2]	R6-58	Page 6-17
[379]	R6-17	Page 6-6	[422]	R6-59	Page 6-17
[380]	R6-18	Page 6-7	[423]	R6-60	Page 6-18
[381]	R6-19	Page 6-7	[424]	R6-61	Page 6-18
[382]	R6-20	Page 6-7	[425]	O6-62	Page 6-18
[383]	R6-21	Page 6-8	[426v2]	R6-63	Page 6-18
[384]	R6-22	Page 6-9	[427]	R6-64	Page 6-18
[385]	R6-23	Page 6-9	[428v2]	R6-65	Page 6-18
[386]	R6-24	Page 6-9	[429]	R6-66	Page 6-18
[387]	R6-25	Page 6-9	[430v3]	R6-70	Page 6-20
[388]	R6-26	Page 6-9	[431v2]	R6-72	Page 6-22
[389]	R6-27	Page 6-9	[432v2]	R6-73	Page 6-22
[390]	R6-28	Page 6-10	[433v2]	R6-78	Page 6-22
[391]	R6-29	Page 6-10	[434]	R6-80	Page 6-22
[392]	R6-30	Page 6-10	[435v2]	R6-81	Page 6-22
[393v2]	R6-31	Page 6-10	[436]	R6-83	Page 6-23
[394]	R6-32	Page 6-10	[437]	R6-84	Page 6-23
[395]	R6-33	Page 6-11	[438]	R6-85	Page 6-23
[397]	R6-34	Page 6-11	[439]	R6-86	Page 6-23
[398]	R6-35	Page 6-11	[440]	R6-87	Page 6-23
[399]	CR6-36	Page 6-11	[441]	R6-88	Page 6-24
[400]	R6-37	Page 6-11	[442]	CR6-89	Page 6-24
[402]	R6-38	Page 6-11	[443]	O6-90	Page 6-24

[444]	R6-91	Page 6-24	[491]	R6-137	Page 6-34
[445]	R6-92	Page 6-24	[492]	O6-138	Page 6-35
[446]	R6-93	Page 6-25	[495]	R6-139	Page 6-35
[447]	R6-94	Page 6-26	[496]	O6-140	Page 6-35
[448]	R6-95	Page 6-26	[497v2]	R6-141	Page 6-35
[449]	R6-96	Page 6-26	[498]	R6-142	Page 6-35
[450]	R6-97	Page 6-26	[499]	R6-143	Page 6-35
[451]	R6-98	Page 6-26	[501]	R6-144	Page 6-35
[452]	R6-99	Page 6-26	[502]	O6-145	Page 6-36
[453]	O6-100	Page 6-27	[505]	R6-146	Page 6-36
[454]	R6-101	Page 6-27	[506]	O6-147	Page 6-36
[455]	R6-102	Page 6-27	[508]	R6-148	Page 6-36
[456]	R6-103	Page 6-27	[509]	R6-149	Page 6-36
[457]	R6-104	Page 6-27	[512v2]	R6-163	Page 6-43
[458]	R6-105	Page 6-27	[513v2]	R6-164	Page 6-44
[459]	R6-106	Page 6-28	[514]	R6-165	Page 6-44
[460]	R6-107	Page 6-28	[515]	R6-166	Page 6-44
[461]	R6-108	Page 6-28	[516]	R6-167	Page 6-44
[462]	R6-109	Page 6-28	[517]	R6-168	Page 6-44
[463]	R6-110	Page 6-28	[518]	R6-169	Page 6-44
[464]	R6-111	Page 6-28	[519v3]	R6-170	Page 6-45
[465]	R6-112	Page 6-29	[521v2]	R6-171	Page 6-45
[466]	R6-113	Page 6-29	[523]	R6-172	Page 6-45
[467v2]	R6-114	Page 6-29	[524]	R6-173	Page 6-46
[468]	R6-115	Page 6-29	[525]	R6-175	Page 6-46
[469]	R6-121	Page 6-30	[526]	R6-176	Page 6-46
[470]	R6-122	Page 6-30	[527]	R6-177	Page 6-46
[471]	R6-123	Page 6-30	[528v4]	R6-178	Page 6-47
[472]	R6-124	Page 6-31	[529]	R6-179	Page 6-47
[473]	O6-125	Page 6-31	[531v2]	R6-180	Page 6-47
[476]	R6-126	Page 6-31	[533]	R6-181	Page 6-47
[477]	O6-127	Page 6-32	[534]	R6-182	Page 6-47
[478v2]	R6-128	Page 6-32	[535]	R6-183	Page 6-48
[479]	R6-129	Page 6-32	[536]	R6-185	Page 6-48
[480v2]	R6-130	Page 6-32	[537]	R6-186	Page 6-48
[481]	R6-131	Page 6-32	[538]	R6-187	Page 6-49
[482]	O6-132	Page 6-32	[539v2]	R6-188	Page 6-49
[485]	R6-133	Page 6-32	[540]	R6-190	Page 6-49
[486]	O6-134	Page 6-32	[541v2]	R6-191	Page 6-49
[488]	R6-135	Page 6-32	[542]	R6-192	Page 6-50
[489v2]	R6-136	Page 6-33	[544v2]	CR6-193	Page 6-50

[547]	R6-196	Page 6-50	[593v2]	CR6-265	Page 6-64
[548]	R6-197	Page 6-51	[595v2]	R6-268	Page 6-65
[549v2]	R6-200	Page 6-51	[596v2]	R6-271	Page 6-65
[550v2]	R6-201	Page 6-52	[597v2]	O6-272	Page 6-65
[551v2]	R6-202	Page 6-52	[598]	R6-273	Page 6-65
[552v2]	R6-204	Page 6-53	[599]	O6-274	Page 6-65
[553v2]	R6-205	Page 6-53	[609v2]	R6-278	Page 6-67
[554]	R6-206	Page 6-53	[610v3]	R6-275	Page 6-67
[555]	R6-207	Page 6-53	[611v2]	R6-279	Page 6-67
[556v2]	R6-209	Page 6-54	[612]	R6-276	Page 6-67
[557v2]	O6-210	Page 6-54	[617v2]	R6-281	Page 6-68
[558]	R6-212	Page 6-54	[618v2]	R6-282	Page 6-68
[559v2]	R6-214	Page 6-55	[620]	R6-291	Page 6-90
[560v2]	O6-215	Page 6-55	[621]	R6-283	Page 6-84
[561]	R6-216	Page 6-55	[622v2]	R6-284	Page 6-84
[562v2]	R6-218	Page 6-55	[623]	R6-292	Page 6-90
[564v2]	R6-220	Page 6-55	[625]	R6-288	Page 6-89
[566v2]	R6-221	Page 6-55	[626v2]	R6-285	Page 6-85
[567v2]	R6-223	Page 6-56	[627v2]	R6-289	Page 6-89
[568v2]	R6-225	Page 6-57	[628v2]	R6-290	Page 6-89
[569v2]	O6-226	Page 6-57	[629v3]	R6-286	Page 6-86
[570]	R6-228	Page 6-58	[632]	R6-293	Page 6-90
[571v2]	R6-230	Page 6-58	[633]	R6-296	Page 6-92
[572v2]	O6-231	Page 6-58	[634v2]	R6-297	Page 6-98
[573]	R6-232	Page 6-58	[635v3]	R6-302	Page 6-99
[574v2]	R6-234	Page 6-59	[636]	R6-299	Page 6-98
[576v2]	R6-236	Page 6-59	[637v2]	R6-300	Page 6-98
[578v2]	R6-237	Page 6-59	[638v3]	R6-304	Page 6-99
[579v2]	R6-239	Page 6-59	[639v2]	R6-298	Page 6-98
[580v3]	R6-240	Page 6-60	[640v3]	R6-308	Page 6-101
[581v2]	R6-242	Page 6-61	[641]	R6-309	Page 6-101
[582v2]	R6-243	Page 6-61	[642]	R6-310	Page 6-101
[583v2]	R6-244	Page 6-61	[643v4]	O6-311	Page 6-106
[584v2]	R6-248	Page 6-62	[645v3]	R6-316	Page 6-107
[585v2]	R6-256	Page 6-63	[656v2]	R6-317	Page 6-108
[586v2]	R6-249	Page 6-62	[657]	R6-318	Page 6-108
[587v2]	R6-258	Page 6-63	[658]	CR6-319	Page 6-108
[588v2]	R6-259	Page 6-63	[659v2]	R6-320	Page 6-108
[589v2]	CR6-261	Page 6-64	[660v2]	R6-321	Page 6-108
[591v2]	CR6-260	Page 6-64	[661]	R6-323	Page 6-112
[592v2]	R6-267	Page 6-65	[662]	R6-324	Page 6-112

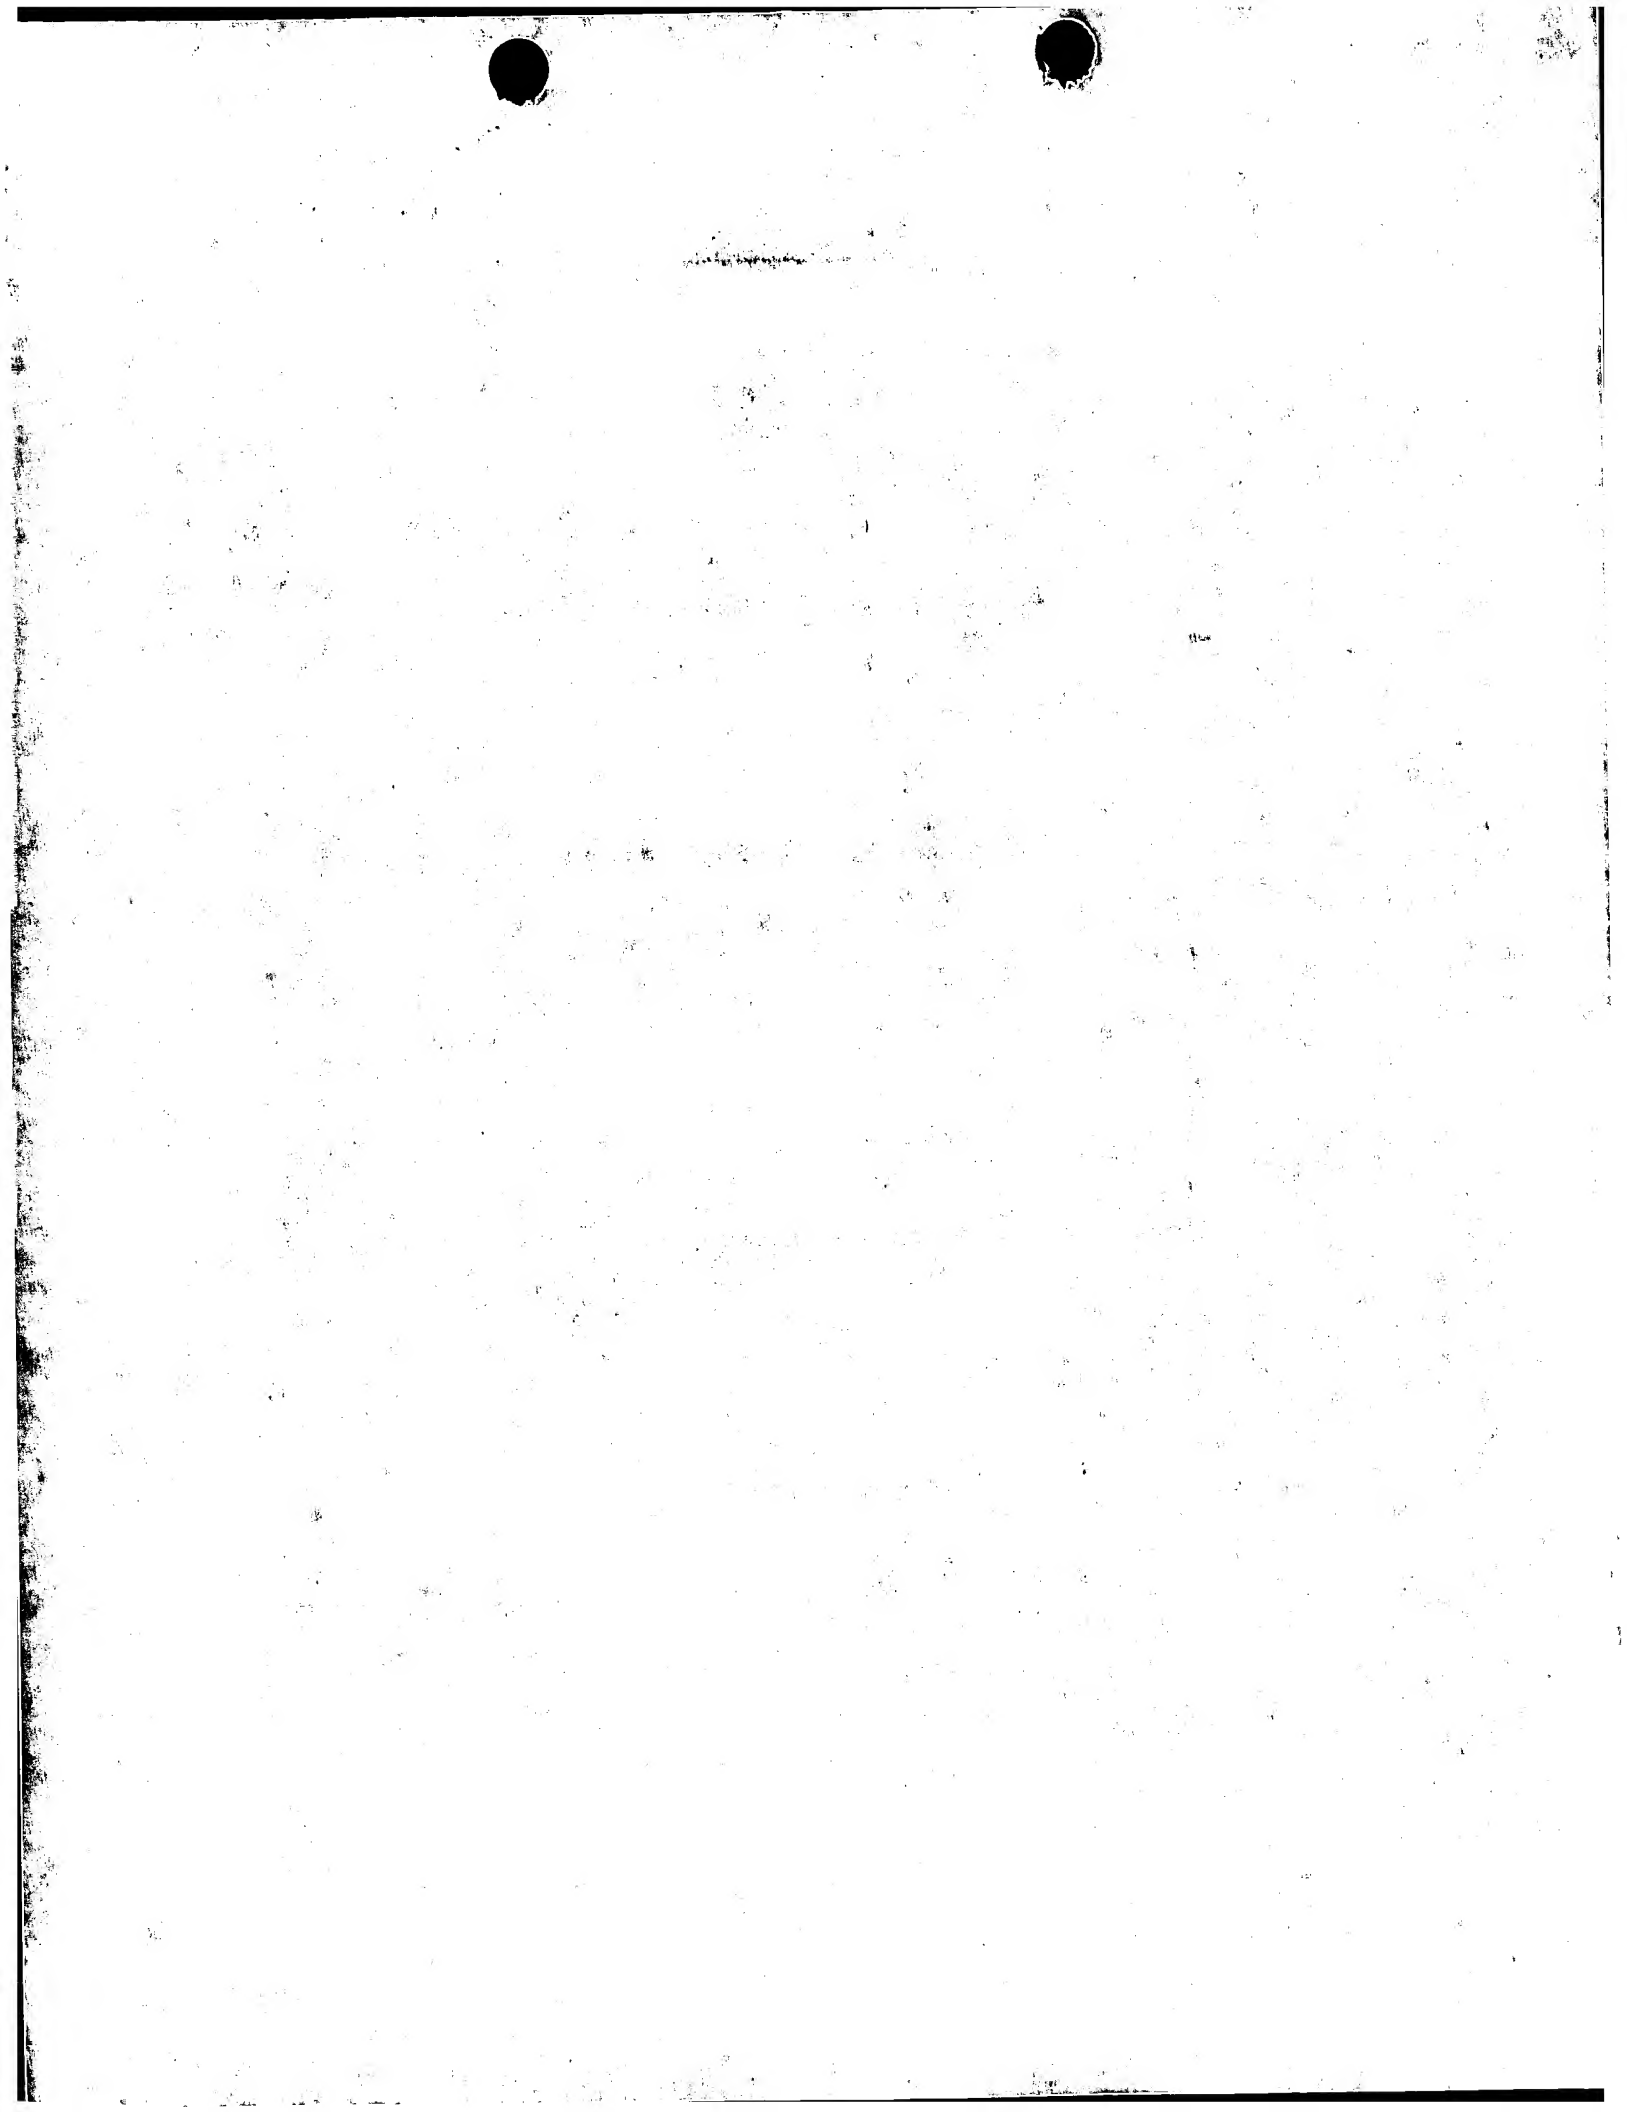
[663]	R6-325	Page 6-112	[719]	R6-378	Page 6-138
[664]	R6-326	Page 6-112	[720]	O6-379	Page 6-138
[667]	R6-327	Page 6-112	[721]	R6-380	Page 6-138
[668]	R6-328	Page 6-112	[722v2]	O6-381	Page 6-138
[669]	R6-329	Page 6-116	[723v2]	O6-382	Page 6-138
[670]	R6-330	Page 6-116	[724v2]	O6-383	Page 6-138
[672]	R6-334	Page 6-116	[725v2]	R6-150	Page 6-38
[673]	R6-335	Page 6-117	[728]	R6-395	Page 6-141
[674]	R6-336	Page 6-120	[729]	CR6-396	Page 6-141
[675]	R6-337	Page 6-120	[730]	O6-397	Page 6-141
[676]	R6-341	Page 6-121	[731v3]	R6-400	Page 6-142
[677]	R6-342	Page 6-121	[732v2]	R6-406	Page 6-142
[678]	R6-343	Page 6-121	[735v3]	R6-404	Page 6-142
[679]	R6-344	Page 6-121	[736v2]	R6-405	Page 6-142
[680]	R6-345	Page 6-121	[737]	R6-410	Page 6-144
[681]	R6-346	Page 6-121	[738]	R6-413	Page 6-144
[682]	R6-349	Page 6-123	[740v2]	R6-387	Page 6-139
[683v3]	R6-347	Page 6-122	[741]	R6-388	Page 6-140
[685v2]	CR6-350	Page 6-132	[743]	R6-389	Page 6-140
[686]	CR6-353	Page 6-132	[744]	R6-391	Page 6-140
[687v2]	R6-354	Page 6-132	[745]	R6-390	Page 6-140
[688v2]	R6-355	Page 6-132	[746]	R6-392	Page 6-140
[689v2]	R6-357	Page 6-132	[747]	R6-393	Page 6-140
[690]	R6-359	Page 6-133	[748v2]	R6-394	Page 6-140
[691v4]	R6-360	Page 6-133	[749]	O6-414	Page 6-146
[693]	O6-362	Page 6-134	[750]	R6-415	Page 6-146
[694v3]	R6-363	Page 6-134	[751v2]	R6-416	Page 6-146
[695v2]	R6-294	Page 6-91	[755]	R6-418	Page 6-147
[703v3]	O6-295	Page 6-91	[756]	R6-417	Page 6-147
[706]	R6-366	Page 6-136	[757v2]	O6-419	Page 6-147
[707]	O6-367	Page 6-136	[759]	R6-421	Page 6-149
[709]	O6-368	Page 6-136	[760]	R6-422	Page 6-149
[710]	R6-369	Page 6-136	[761]	O7-1	Page 7-2
[711]	R6-370	Page 6-136	[762]	R7-2	Page 7-2
[712]	R6-371	Page 6-136	[763]	R7-3	Page 7-2
[713]	R6-372	Page 6-136	[764]	R7-4	Page 7-3
[714]	R6-373	Page 6-137	[765]	R7-5	Page 7-3
[715]	CR6-374	Page 6-137	[766]	R7-6	Page 7-3
[716]	R6-375	Page 6-137	[767]	R7-7	Page 7-3
[717]	R6-376	Page 6-137	[768]	R7-8	Page 7-3
[718]	R6-377	Page 6-137	[771]	R7-9	Page 7-3

[772]	R7-10	Page 7-3	[813v3]	R8-41	Page 8-16
[773]	R7-11	Page 7-4	[814]	R8-43	Page 8-17
[774]	R7-12	Page 7-4	[815]	R8-44	Page 8-17
[775]	O7-13	Page 7-5	[816]	R8-45	Page 8-17
[776]	R8-1	Page 8-7	[817]	R8-46	Page 8-17
[777]	R8-2	Page 8-7	[818]	CR8-47	Page 8-17
[778]	R8-3	Page 8-7	[819v2]	R8-48	Page 8-18
[779]	O8-4	Page 8-7	[820]	O8-49	Page 8-18
[780]	R8-5	Page 8-8	[821]	O8-50	Page 8-18
[781]	R8-6	Page 8-8	[822]	R8-51	Page 8-18
[782]	R8-7	Page 8-8	[823]	R8-52	Page 8-20
[783]	R8-8	Page 8-9	[824]	R8-53	Page 8-20
[784]	R8-9	Page 8-10	[825]	R8-54	Page 8-20
[785]	R8-11	Page 8-10	[826]	R8-55	Page 8-21
[786]	R8-12	Page 8-11	[827]	R8-56	Page 8-21
[787]	R8-13	Page 8-11	[828]	R8-57	Page 8-21
[788]	R8-14	Page 8-11	[829]	R8-58	Page 8-21
[789]	R8-15	Page 8-12	[830]	R8-59	Page 8-21
[790]	R8-16	Page 8-12	[831]	R8-60	Page 8-21
[791]	R8-17	Page 8-12	[832]	R8-61	Page 8-21
[792]	CR8-18	Page 8-12	[833]	R8-62	Page 8-21
[793v2]	CR8-19	Page 8-12	[834]	R8-63	Page 8-21
[794v2]	R8-21	Page 8-13	[835]	O8-64	Page 8-21
[795v2]	R8-22	Page 8-13	[836]	R8-65	Page 8-22
[796]	R8-24	Page 8-13	[837]	R8-66	Page 8-22
[797]	R8-25	Page 8-14	[838]	R8-67	Page 8-22
[798]	R8-26	Page 8-14	[839]	R8-69	Page 8-22
[799v2]	R8-27	Page 8-14	[841]	R8-70	Page 8-22
[800]	R8-28	Page 8-15	[842]	R8-71	Page 8-22
[801]	R8-29	Page 8-15	[843]	R8-72	Page 8-22
[802]	R8-30	Page 8-15	[844]	R8-73	Page 8-23
[803]	R8-31	Page 8-15	[845]	R8-74	Page 8-23
[804]	R8-32	Page 8-15	[846]	R8-75	Page 8-23
[805]	R8-33	Page 8-15	[847]	R8-76	Page 8-24
[806]	R8-34	Page 8-16	[848]	R8-77	Page 8-24
[807]	R8-35	Page 8-16	[849]	R8-78	Page 8-24
[808]	R8-36	Page 8-16	[850v2]	R8-79	Page 8-24
[809]	R8-37	Page 8-16	[851]	R8-80	Page 8-25
[810]	R8-38	Page 8-16	[852]	R8-81	Page 8-25
[811]	R8-39	Page 8-16	[853]	R8-82	Page 8-25
[812]	R8-40	Page 8-16	[854]	R8-84	Page 8-25

[855]	R8-85	Page 8-27	[896v2]	R8-132	Page 8-50
[856]	R8-86	Page 8-27	[897]	R8-133	Page 8-50
[857]	R8-87	Page 8-27	[898]	R8-134	Page 8-50
[858]	R8-88	Page 8-28	[899]	R8-135	Page 8-51
[859]	R8-89	Page 8-28	[900v2]	R8-136	Page 8-51
[860]	R8-90	Page 8-28	[901]	R8-137	Page 8-51
[861]	R8-91	Page 8-29	[902]	R8-138	Page 8-51
[862]	R8-92	Page 8-29	[903]	R3-54	Page 3-50
[863]	R8-93	Page 8-29	[904]	R3-61	Page 3-51
[864]	R8-94	Page 8-29	[905]	R3-68	Page 3-53
[865]	R8-95	Page 8-29	[906]	R3-72	Page 3-55
[866]	R8-96	Page 8-29	[907]	R3-73	Page 3-55
[867]	R8-97	Page 8-30	[908]	O3-101	Page 3-74
[868]	R8-98	Page 8-30	[909]	O3-114	Page 3-80
[869]	R8-99	Page 8-36	[910]	CR4-20	Page 4-43
[870]	O8-100	Page 8-36	[911]	R5-4	Page 5-2
[871]	O8-101	Page 8-36	[912]	R5-5	Page 5-2
[872]	O8-102	Page 8-36	[913]	CR5-49	Page 5-24
[873]	R8-103	Page 8-36	[914]	R5-54	Page 5-26
[874]	O8-104	Page 8-36	[915]	CR5-55	Page 5-27
[875]	R8-105	Page 8-37	[916]	R5-56	Page 5-27
[876]	R8-106	Page 8-38	[917]	O5-86	Page 5-40
[877]	R8-107	Page 8-38	[918]	R5-98	Page 5-52
[878v2]	R8-108	Page 8-38	[919]	R5-108	Page 5-57
[879]	R8-109	Page 8-38	[920]	R5-109	Page 5-58
[880v2]	R8-110	Page 8-38	[921]	R5-120	Page 5-61
[881v2]	R8-111	Page 8-39	[922]	R5-123	Page 5-62
[882]	R8-114	Page 8-39	[923]	R5-124	Page 5-62
[883v2]	R8-116	Page 8-40	[924]	R5-125	Page 5-62
[884]	R8-117	Page 8-40	[925]	R5-126	Page 5-62
[885]	R8-118	Page 8-40	[926]	R5-130	Page 5-64
[886]	R8-119	Page 8-40	[928]	O5-139	Page 5-72
[887]	R8-121	Page 8-43	[930v2]	R5-146	Page 5-73
[888]	R8-122	Page 8-43	[931]	R5-233	Page 5-105
[889]	R8-123	Page 8-43	[932]	R5-249	Page 5-115
[890]	R8-124	Page 8-45	[933]	R5-250	Page 5-115
[891v2]	R8-125	Page 8-45	[934]	R5-251	Page 5-115
[892]	CR8-126	Page 8-45	[935]	R5-252	Page 5-115
[893]	R8-128	Page 8-49	[936]	O5-253	Page 5-115
[894]	R8-130	Page 8-49	[937v2]	O5-261	Page 5-127
[895]	R8-131	Page 8-50	[938]	R6-3	Page 6-2

[939]	R6-4	Page 6-4	[980]	R6-262	Page 6-64
[940]	O6-53	Page 6-16	[981]	R6-263	Page 6-64
[941]	R6-67	Page 6-19	[982]	R6-264	Page 6-64
[942]	CR6-68	Page 6-19	[983]	R6-266	Page 6-65
[943]	CR6-69	Page 6-19	[984]	CR6-269	Page 6-65
[944v2]	R6-71	Page 6-20	[985]	CR6-270	Page 6-65
[945]	R6-74	Page 6-22	[986]	R6-277	Page 6-67
[946]	R6-75	Page 6-22	[987]	R6-280	Page 6-68
[947]	R6-76	Page 6-22	[988]	R6-287	Page 6-88
[948]	R6-77	Page 6-22	[989]	R6-301	Page 6-98
[949]	R6-79	Page 6-22	[990]	R6-303	Page 6-99
[950]	R6-82	Page 6-22	[991v2]	R6-306	Page 6-100
[951]	R6-189	Page 6-49	[992v2]	R6-307	Page 6-100
[952]	CR6-194	Page 6-50	[994v3]	O6-312	Page 6-106
[953]	O6-195	Page 6-50	[995v2]	R6-348	Page 6-123
[954]	R6-198	Page 6-51	[996]	R6-365	Page 6-134
[955]	R6-199	Page 6-51	[997v2]	R6-151	Page 6-38
[956]	O6-203	Page 6-52	[999v2]	CR6-398	Page 6-141
[957v2]	O6-208	Page 6-54	[1000v2]	R6-399	Page 6-141
[958]	R6-211	Page 6-54	[1001v2]	O6-157	Page 6-41
[959]	O6-213	Page 6-54	[1002v2]	CR6-401	Page 6-142
[960]	R6-217	Page 6-55	[1003]	R6-402	Page 6-142
[961]	R6-219	Page 6-55	[1004]	O6-403	Page 6-142
[962v2]	R6-222	Page 6-56	[1005v2]	R6-407	Page 6-143
[963v2]	O6-224	Page 6-57	[1006v2]	R6-408	Page 6-143
[964]	R6-227	Page 6-57	[1007v2]	CR6-409	Page 6-143
[965]	O6-229	Page 6-58	[1008]	R6-411	Page 6-144
[966]	R6-233	Page 6-58	[1009]	CR6-412	Page 6-144
[967]	R6-235	Page 6-59	[1010]	O6-420	Page 6-147
[968v2]	R6-238	Page 6-59	[1011]	CR8-23	Page 8-13
[969]	R6-241	Page 6-60	[1012]	R8-129	Page 8-49
[970]	CR6-245	Page 6-62	[1013]	R5-8	Page 5-11
[971]	R6-246	Page 6-62	[1014]	R5-137	Page 5-70
[972]	R6-247	Page 6-62	[1015]	R5-149	Page 5-73
[973]	R6-250	Page 6-62	[1016]	CR5-150	Page 5-73
[974]	R6-251	Page 6-62	[1017]	R5-151	Page 5-73
[975]	R6-252	Page 6-62	[1018]	CR5-154	Page 5-76
[976]	CR6-253	Page 6-63	[1019]	R5-155	Page 5-76
[977]	R6-254	Page 6-63	[1020]	R5-156	Page 5-76
[978]	R6-255	Page 6-63	[1021]	R5-157	Page 5-77
[979]	R6-257	Page 6-63	[1022]	R5-191	Page 5-87

[1023]	O5-196	Page 5-90	[1064]	R6-116	Page 6-29
[1024]	R5-213	Page 5-95	[1065]	R6-117	Page 6-29
[1025]	R5-239	Page 5-110	[1066]	R6-118	Page 6-29
[1026]	R5-240	Page 5-110	[1067]	R6-119	Page 6-30
[1027]	R5-255	Page 5-120	[1068]	R6-120	Page 6-30
[1028]	R5-257	Page 5-122	[1069]	R6-152	Page 6-39
[1029]	R5-258	Page 5-124	[1070]	R6-153	Page 6-39
[1030]	R5-259	Page 5-125	[1071]	R6-154	Page 6-40
[1031v2]	R5-262	Page 5-127	[1072]	R6-155	Page 6-40
[1032v2]	O5-263	Page 5-127	[1073]	R6-156	Page 6-41
[1033]	O6-322	Page 6-109	[1074]	O6-158	Page 6-41
[1034v2]	R6-364	Page 6-134	[1075]	R6-159	Page 6-41
[1035]	O8-10	Page 8-10	[1076]	O6-160	Page 6-41
[1036]	O8-20	Page 8-13	[1077]	R6-161	Page 6-41
[1037]	R8-68	Page 8-22	[1078]	R6-162	Page 6-42
[1038]	R8-83	Page 8-25	[1079]	O6-174	Page 6-46
[1039]	O8-112	Page 8-39	[1080]	O6-184	Page 6-48
[1040v2]	R8-113	Page 8-39	[1081]	R6-305	Page 6-99
[1041]	R8-115	Page 8-39	[1082]	R6-313	Page 6-106
[1042]	R8-120	Page 8-42	[1083]	O6-314	Page 6-106
[1043]	R8-127	Page 8-46	[1084]	O6-315	Page 6-106
[1044]	R3-75	Page 3-57	[1085]	CR6-331	Page 6-116
[1045]	O3-76	Page 3-58	[1086]	R6-332	Page 6-116
[1046]	R5-51	Page 5-25	[1087]	R6-333	Page 6-116
[1047]	R5-52	Page 5-25	[1088]	CR6-338	Page 6-120
[1048]	CR5-101	Page 5-54	[1089]	R6-339	Page 6-121
[1049]	CR5-102	Page 5-54	[1090]	R6-340	Page 6-121
[1050]	R5-103	Page 5-54	[1091]	CR6-351	Page 6-132
[1051]	R5-110	Page 5-58	[1092]	O6-352	Page 6-132
[1052]	CR5-158	Page 5-77	[1093]	R6-356	Page 6-132
[1053]	R5-159	Page 5-77	[1094]	O6-358	Page 6-133
[1054]	CR5-184	Page 5-86	[1095]	R6-361	Page 6-133
[1055]	R5-185	Page 5-86	[1096]	O6-384	Page 6-138
[1056]	R5-186	Page 5-86	[1097]	O6-385	Page 6-139
[1057]	CR5-187	Page 5-86	[1098]	O6-386	Page 6-139
[1058]	R5-188	Page 5-86	[1099]	R8-42	Page 8-17
[1059]	R5-189	Page 5-87			
[1060]	R5-190	Page 5-87			
[1061]	CR5-214	Page 5-95			
[1062]	R5-215	Page 5-95			
[1063]	R5-216	Page 5-95			



**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☒ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☒ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☒ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE IS BLANK